

UNIVERSAL GRAPHS WITH FORBIDDEN SUBGRAPHS AND ALGEBRAIC CLOSURE

Gregory Cherlin ¹

Dept. Mathematics, Rutgers University, Busch Campus, New Brunswick, NJ 08903

Saharon Shelah ²

Dept. Mathematics, Hebrew University, Jerusalem, Israel, and Rutgers University, New Brunswick, NJ 08903

Niandong Shi

Dept. Mathematics, East Stroudsburg University, East Stroudsburg, PA 18301

ABSTRACT

We apply model theoretic methods to the problem of existence of countable universal graphs with finitely many forbidden connected subgraphs. We show that to a large extent the question reduces to one of local finiteness of an associated “algebraic closure” operator (Theorem 3, §3). The main applications are new examples of universal graphs with forbidden subgraphs (§§7, 8, and 10) and simplified treatments of some previously known cases (§§6.2, 6.3).

Introduction

We are concerned here with the following problem: give a finite set \mathcal{C} of finite connected graphs, does the class $\mathcal{G}_{\mathcal{C}}$ of countable graphs which omit \mathcal{C} contain a universal element (one in which all others are embeddable as induced subgraphs)? Here we say that a graph G *omits* a class \mathcal{C} of graphs if no graph in \mathcal{C} embeds as a subgraph of G . The problem is to characterize those classes \mathcal{C} for which there is such a universal graph. A more fundamental problem is whether there is any effective characterization of these classes \mathcal{C} ; in other words, is there an algorithm which will produce the answer in each case? This problem remains open even when \mathcal{C} consists of a *single* forbidden subgraph, though an accumulation of evidence, some given in the present paper, suggests that at least this instance should have an explicit and fairly simple solution. We discuss this further below.

Rado observed [Ra] that there is a universal countable graph. This corresponds to the case $\mathcal{C} = \emptyset$. Many other cases have been considered in the literature [ChK, CS1, CS2, CST, FK1, Ko, FK2, GK, KMP, Ko, KP1, KP2, Pa]. In particular [FK1] gives a complete solution for the case in which \mathcal{C} consists of a single 2-connected constraint, and [CST] treats the case in which \mathcal{C} consists of a single tree with no vertex of degree 2.

To date very few cases have been identified in which a universal countable \mathcal{C} -free graph exists. For $\mathcal{C} = \{C\}$ consisting of a single constraint, the following cases are known to allow a universal graph: C is complete; C is a tree consisting of one path to which at most one additional edge is attached; or C is a “bow-tie”, a particular graph of order 5. We will add some additional families of examples using model theoretic methods.

Another family of universal \mathcal{C} -free graphs corresponds to the class \mathcal{C} of odd cycles of length up to some specified bound (C_{2n+1} -free graphs for $n \leq N$). We generalize this to the case in which \mathcal{C} is closed under homomorphism in an appropriate sense (Theorem 4, and §7). In earlier work the positive results have generally come either from structure theorems for \mathcal{C} -free graphs (notably in the path-free case [KMP]) or

¹ Research supported in part by NSF grant DMS 9501176

² Research supported in part by U.S.-Israel Binational Science Foundation grant 0377215

from Fraïssé’s amalgamation method, whose complexity increases rapidly as the constraint class \mathcal{C} becomes more complicated. Using the model theoretic notions of existential completeness and algebraic closure for \mathcal{C} -free graphs, we give a criterion for the existence of a universal \mathcal{C} -free graph which effectively short circuits this process. Our arguments can be converted into amalgamation arguments in principle, but not in any very explicit way.

Kojman conjectured in conversation years ago that closure of the constraint class under homomorphic image might be a key condition in connection with the existence of universal graphs. Our work confirms this conjecture in one direction (Theorems 4 and 5) and relates the condition directly to the broader issue of the structure of the algebraic closure operator.

Our model theoretic methods are very close to those which have been used in practice in all cases in which *nonexistence* of universal graphs has been established. There they are typically referred to as “rigidity” arguments. This amounts to a rephrasing in purely graph theoretic terms of a more general model theoretic notion. We did this ourselves in [CST], though in fact an awareness of the model theoretic framework lay in the background of the proof given there. In our present work, we have reached the point at which such a reformulation of our methods would be counterproductive, as we make use of general considerations which are well known in model theory but have not yet played in explicit role in graph theory. Accordingly, the first half of the present paper lays the foundation of our approach, recalling what we need from model theory and applying it in the case of \mathcal{C} -free graphs. As we will see in Theorem 3 of §3, these ideas produce much clearer results in the \mathcal{C} -free context than one would get in a more general model theoretic context. This is really the key to our whole analysis.

Applications of these general ideas are found in §§5 – 10. Most of the cases considered in §§5, 6, 9 were treated successfully in the past, and are reexamined from our present point of view partly by way of illustration and partly because our present viewpoint suggests quantitative issues extending the earlier purely qualitative analysis. That is, in cases in which our “algebraic closure” operator is locally finite, we consider its rate of growth.

New examples are given in §§7, 8. In particular §8 is devoted to an infinite family extending the “bow-tie” example considered by Komjáth [Ko] using methods that have further potential. This is the hardest case treated here.

To conclude this introduction we take note of two directions which are particularly promising for further work: the general problem of effectivity, and the case of a single constraint.

Effectivity *Given a finite set \mathcal{C} of finite connected graphs, determine whether there is a universal countable \mathcal{C} -free graph.*

It is by no means clear that this problem should have an effective solution. It is natural to consider a further generalization in which graphs are replaced by vertex-colored, edge-colored, and directed graphs, or more generally by relational structures for any finite relational language. It seems likely however that this more general problem can be reduced to the special case of graphs by a suitable encoding. This is one reason why the existence of an effective solution is doubtful, but at the present time the question is entirely open.

Single constraints *Let \mathcal{U} be the collection of all finite connected graphs C for which there is a countable universal C -free graph, and let*

$$\mathcal{U}_0 = \{C \in \mathcal{U} : \text{every induced subgraph of } C \text{ is in } \mathcal{U}\}.$$

Conjecture $\mathcal{U} = \mathcal{U}_0$.

While it is not likely that this conjecture will be proved *a priori*, it may well turn out to be the case. The point of the conjecture is that it should be possible to determine \mathcal{U}_0 explicitly using known methods, and then rephrase the conjecture more explicitly. In §8 below we will give a new family of examples in \mathcal{U}_0 . What is needed, apparently, is to continue that analysis, which will involve substantial computation, and also to prove a number of further results on nonexistence of universal graphs. Note that by [FK1] any block (maximal 2-connected subgraph) of a graph in \mathcal{U}_0 is complete, and the results of [GK] can be combined with some similar constructions to reduce the class of candidates for members of \mathcal{U}_0 to graphs much like those considered in §8.

I. General theory.

In this part we will discuss the application of some model theoretic ideas to the general problem of the existence of universal countable graphs with forbidden subgraphs. In §1 we associate with a class \mathcal{C} of finite graphs (usually taken to be connected) the class $\mathcal{G}_{\mathcal{C}}$ of countable graphs “omitting” \mathcal{C} and the subclass $\mathcal{E}_{\mathcal{C}}$ of “existentially complete” graphs in $\mathcal{G}_{\mathcal{C}}$. The key to the model theoretic approach is to understand $\mathcal{E}_{\mathcal{C}}$. In fact where a universal \mathcal{C} -free graph exists, it is often the case that $\mathcal{E}_{\mathcal{C}}$ contains a unique graph, up to isomorphism, and this graph is then a “canonical” universal \mathcal{C} -free graph. Using well established model theoretic terminology, we refer to this as the \aleph_0 -categorical case. The role of $\mathcal{E}_{\mathcal{C}}$ in connection with the problem of determining whether a universal \mathcal{C} -free graph exists is explored in §2, which introduces the important technical notion of an *existential type*. In §3 we characterize the \aleph_0 -categorical case in terms of the behavior of the associated algebraic closure operator on $\mathcal{E}_{\mathcal{C}}$. We begin the analysis of the algebraic closure operator in §4. More delicate techniques for analyzing this operator are left to the second part, as needed for applications.

Our thanks go to P. Komjath for a close reading of a draft of the present paper.

§1. Existentially complete \mathcal{C} -free graphs.

First we introduce some definitions and notations which will be used in the whole paper.

Definition 1 *Let \mathcal{C} be a set of finite graphs.*

1. *A graph G omits \mathcal{C} if no subgraph of G is isomorphic to any graph in \mathcal{C} .*
2. *$\mathcal{G}_{\mathcal{C}}$ is the class of all countable graphs omitting \mathcal{C} .*
3. *A graph $G \in \mathcal{G}_{\mathcal{C}}$ is universal (for $\mathcal{G}_{\mathcal{C}}$) if every graph in $\mathcal{G}_{\mathcal{C}}$ is isomorphic to an induced subgraph of G .*

Remarks

1. There are two notions of universality which are generally considered. We say that $G \in \mathcal{G}_{\mathcal{C}}$ is *weakly universal* if every graph in $\mathcal{G}_{\mathcal{C}}$ is isomorphic to a subgraph of G . In practice the two notions of universality behave similarly. A universal graph is evidently weakly universal, and in practice proofs of the nonexistence of a universal graph can often be doctored in standard ways to exclude weakly universal graphs as well.

For a theoretical analysis our definition of universality is to be preferred, at least initially, as it facilitates the application of general methods. To pass to the weakly universal case on a theoretical level is in part a

matter of replacing “existential type” in §2 by “positive existential types”, but the more pragmatic alternative of working mainly with (strictly) universal graphs on a theoretical level and then doctoring specific construction is probably to be preferred.

2. Let $T_{\mathcal{C}}$ be the *first order theory* of $\mathcal{G}_{\mathcal{C}}$. Then the models of $T_{\mathcal{C}}$ are all the \mathcal{C} -free graphs and $\mathcal{G}_{\mathcal{C}}$ consists of the countable models of $T_{\mathcal{C}}$, which is a *universal theory*. This reflects the assumption that all graphs in \mathcal{C} are finite, and allows the application of model theoretic methods.

Definition 2 *Let \mathcal{C} be a set of finite graphs.*

1. If $G \subseteq H$ are graphs, we say that G is *existentially complete in H* if every existential statement ϕ which is defined in G and true in H is also true in G . Equivalently, if $A \subseteq B$ are finite induced subgraphs of G and H respectively, then there is an embedding $f : B \rightarrow G$ taking B isomorphically onto an induced subgraph of G , with $f \upharpoonright A$ the identity.

2. $G \in \mathcal{G}_{\mathcal{C}}$ is said to be *existentially complete (for $\mathcal{G}_{\mathcal{C}}$)* if G is existentially complete in each graph H such that $G \subseteq H \in \mathcal{G}_{\mathcal{C}}$.

3. $\mathcal{E}_{\mathcal{C}}$ is the class of all existentially complete graphs in $\mathcal{G}_{\mathcal{C}}$.

4. $T_{\mathcal{C}}^*$ is the theory of $\mathcal{E}_{\mathcal{C}}$. (In the proof of Theorem 1 below we will determine this theory fairly precisely.)

Example 1 If $\mathcal{C} = \emptyset$, then $\mathcal{G}_{\mathcal{C}}$ is the class of all countable graphs and $\mathcal{E}_{\mathcal{C}}$ contains only one element up to isomorphism: the random countable graph G_{∞} [Ra]. $T_{\mathcal{C}}$ is the theory of graphs, and $T_{\mathcal{C}}^*$ is the theory of G_{∞} (a complete theory).

Example 2 If $\mathcal{C} = \{K_3\}$, a complete graph, then $\mathcal{G}_{\mathcal{C}}$ is the class of countable triangle-free graphs and $\mathcal{E}_{\mathcal{C}}$ contains a unique element up to isomorphism, called the generic triangle-free graph G_3 . $T_{\mathcal{C}}$ is the theory of triangle-free graphs and $T_{\mathcal{C}}^*$ is the theory of G_3 . For $\mathcal{C} = \{K_n\}$, any n , the situation is similar.

Example 3 If $\mathcal{C} = \{K_2 + K_2\}$, the disjoint sum of two copies of K_2 , then $\mathcal{E}_{\mathcal{C}}$ contains two elements up to isomorphism: the triangle K_3 and the star S_{∞} of infinite degree. The theory $T_{\mathcal{C}}^*$ is not a complete theory, since K_3 and S_{∞} have different theories.

Example 4 If $\mathcal{C} = \{S_3\}$ (S_n denotes a star of degree n or order $n + 1$), then $T_{\mathcal{C}}$ is the theory of graphs G with vertex degree at most 2, and $T_{\mathcal{C}}^*$ is the theory of graphs in which every vertex has degree 2, and which contain infinitely many cycles C_n for each $n \geq 3$. The countable models G of $T_{\mathcal{C}}^*$ are characterized up to isomorphism by the number of connected components in G isomorphic to a 2-way infinite path. If G_k is the model of $T_{\mathcal{C}}^*$ with k components of this form ($k \geq 0$), then G_{∞} is universal for this class.

Remarks

1. We will see below that $T_{\mathcal{C}}^*$ is complete if the graphs in \mathcal{C} are connected. This is the case of primary interest here.

2. It is easy to see that there is a universal graph in $\mathcal{G}_{\mathcal{C}}$ if and only if there is a universal graph in $\mathcal{E}_{\mathcal{C}}$. We will attempt to make this observation more useful by analyzing $T_{\mathcal{C}}^*$ and $\mathcal{E}_{\mathcal{C}}$ more clearly below.

3. The notion of existential completeness makes sense in almost any context (though our reformulation in terms of embeddings is not always accurate). For example, algebraically closed fields are existentially complete by Hilbert’s Nullstellensatz; real closed fields are existentially complete in the category of ordered fields (Tarski); and dense linear orders are existentially complete in the category of linear orders.

4. While it is natural to think of existential completeness as a form of “algebraic closure”, it does not involve the sort of *finiteness assumptions* connected intuitively with notion of algebraicity. We will introduce the model theoretic algebraic closure operator below.

Theorem 1 *Let \mathcal{C} be a finite set of finite graphs. Then*

1. $\mathcal{E}_{\mathcal{C}}$ is the class of countable models of the theory $T_{\mathcal{C}}^*$.
2. If every $C \in \mathcal{C}$ is connected, then $T_{\mathcal{C}}^*$ is complete.

The proof will involve the general theory of model complete theories and existentially complete structures, as presented in [HW]. We first give an example showing the necessity of taking \mathcal{C} finite.

Example 5 *Let $\mathcal{C} = \{C_n : n \geq 3\}$, the class of all cycles. Then $\mathcal{G}_{\mathcal{C}}$ is the class of countable forests and $\mathcal{E}_{\mathcal{C}}$ contains a unique graph T_{∞} , up to isomorphism, a tree in which every vertex is of countable infinite degree. The models of $T_{\mathcal{C}}^*$ are disjoint unions of any number of copies of T_{∞} .*

Remark *In Theorem 1, clause (2) follows from clause (1). This is because clause (1) provides one of the standard criteria for the theory $T_{\mathcal{C}}^*$ to be model complete (Robinson’s Test, [HW, Theorem 2.2]) and for such theories, completeness is equivalent to the joint embedding property: any two models of a theory should be contained as induced subgraphs in a third [HW, Proposition 2.8]. If \mathcal{C} consists of connected graphs, then $T_{\mathcal{C}}$ is closed under the formation of disjoint sums. However connectedness is not a necessary condition for joint embedding:*

Example 6 *If $\mathcal{C} = \{K_3, K_2 + K_2\}$ then $T_{\mathcal{C}}$ has the joint embedding property.*

It is not clear whether one can easily recognize the finite sets \mathcal{C} for which $T_{\mathcal{C}}$ has the joint embedding property.

The proof of Theorem 1 requires the following technical lemma. Recall that a quantifier-free formula is *conjunctive* if it is a conjunction of atomic formulas and the negations of atomic formulas. An existential formula of the form $\exists \bar{x}\phi$ with ϕ quantifier-free and conjunctive is called *primitive*. A typical example of a conjunctive formula is a description of the isomorphism type of an induced subgraph.

Notation.

Let ϕ be a formula. We write $T_{\mathcal{C}} \vdash \phi$ if every \mathcal{C} -free graph satisfies “ $\forall \bar{x}\phi(\bar{x})$ ” (we quantify over all free variables in ϕ). In other words, ϕ is “always” true in \mathcal{C} -free graphs.

Lemma 1 *Let \mathcal{C} be a finite set of forbidden substructures. For each $n \geq 0$ there is a natural number b_n such that for any two primitive existential formulas ϕ, ψ such that*

- i. ϕ contains at most n existential quantifiers,
 - ii. $T_{\mathcal{C}} \vdash \neg(\phi \wedge \psi)$, and
 - iii. for each pair of variables y_1, y_2 occurring in ψ , with at least one of them quantified, the clause $y_1 \neq y_2$ occurs as a conjunct in ψ ,
- there is a subformula ψ_1 of ψ such that*

1. ψ_1 contains at most b_n existential quantifiers.
2. $T_{\mathcal{C}} \vdash \neg(\phi \wedge \psi_1)$.

We will first explain how Theorem 1 follows from this lemma, then prove the lemma. The following is essentially a corollary to Lemma 1.

Lemma 2 *Let $\phi(\bar{x})$ be a universal formula. Then there is an existential formula $\psi(\bar{x})$ such that*

$$T_C \vdash \forall \bar{x}[\phi(\bar{x}) \longleftrightarrow \psi(\bar{x})].$$

Proof:

We use Proposition 1.6 (iii) of [HW]. Let Φ be the set of all existential formulas $\phi'(\bar{x})$ such that

$$T_C \vdash \forall \bar{x}[\phi'(\bar{x}) \longrightarrow \phi(\bar{x})].$$

Then for $G \in \mathcal{E}_C$, $\bar{u} \in G$, we have

$$G \models \phi(\bar{u}) \iff G \models \phi'(\bar{u}) \text{ for some } \phi' \in \Phi.$$

In other words,

$$G \models \forall \bar{x}[\phi(\bar{x}) \longleftrightarrow \bigvee_{\Phi} \phi'(\bar{x})]. \quad (*)$$

Note that the disjunction on the right is infinite; using Lemma 1 we will replace Φ by a finite subset Φ' for which the analog of $(*)$ holds. Thus with $\psi = \bigvee_{\Phi'} \phi'$, the claim follows.

Any existential formula is equivalent to a disjunction of primitive existential formulas; so we may take Φ to consist of primitive existential formulas. Similarly, the universal formula ϕ is equivalent to a conjunction of negations of primitive existential formulas, so it suffices to deal with the case $\phi = \neg\phi_1$ with ϕ_1 primitive existential. Finally, we may suppose that for each $\phi' \in \Phi$ and each pair y_1, y_2 of variables occurring existentially quantified in ϕ' , we have $y_i \neq y_j$ as a conjunct in ϕ' for $i \neq j$. Indeed if $\phi' = \exists \bar{y} \phi'_0(\bar{x}, \bar{y})$, then

$$\phi' \longleftrightarrow \exists \bar{y}(\phi'_0 \wedge y_i = y_j) \vee \exists \bar{y}(\phi'_0 \wedge y_i \neq y_j)$$

so we may replace ϕ' if necessary by two disjuncts on the right and then contract variables in the first disjunct.

After these preparations, Φ consists of formulas ϕ' to which Lemma 1 applies, with n the number of quantifiers occurring in ϕ . Thus if $\Phi' \subseteq \Phi$ consists of the primitive existential formulas ϕ' , in at most b_n variables such that

$$T_C \models \neg(\phi_1 \wedge \phi'),$$

then

$$T_C \models \forall \bar{x}[\phi(\bar{x}) \longleftrightarrow \bigvee_{\Phi'} \phi'(\bar{x})].$$

■

Proof of Theorem 1:

By Lemma 2, every universal formula is equivalent to an existential formula modulo T_C . This is equivalent to clause (1) of Theorem 1 by [CK, Theorem 3.5.1]. As noted before, clause (2) follows from clause (1). ■

Proof of Lemma 1:

We proceed by induction on n , the number of quantified variables in $\phi(\bar{x})$. Let $k = \max\{|C| : C \in \mathcal{C}\}$.

If $n = 0$, then ϕ is quantifier free and we will take $b_0 = k$. Suppose $T_C \models \neg(\phi \wedge \psi)$. We have $\psi = \exists \bar{y} \psi_0(\bar{x}, \bar{y})$, ψ_0 either contradicts ϕ explicitly, or states that the induced graph on some subset of k vertices from \bar{x}, \bar{y} contains a forbidden subgraph. In the former case ψ can be replaced by a quantifier free formula, and in the latter case by a formula in at most k quantified variables.

For the induction step, let $\phi = \exists \bar{y} \phi_0(\bar{x}, \bar{y})$ have $n + 1$ quantified variables and let $\psi = \exists \bar{y}' \psi_0(\bar{x}, \bar{y}')$. Let A and B be the graphs on vertices \bar{x}, \bar{y}' described by ϕ_0 and ψ_0 respectively, that is, edges exist as specified by ϕ_0 and ψ_0 . As $T_C \models \neg(\phi \wedge \psi)$ the free joint of A and B over \bar{x} contains a forbidden graph $C \in \mathcal{C}$. For each pair of variables y_i in $C \cap \bar{y}$ and y'_j in $C \cap \bar{y}'$, introduce a new variable x_{ij} and let $\phi_0^*(\bar{x}, x_{ij}, \hat{y})$ and $\psi_0^*(\bar{x}, x_{ij}, \hat{y}')$ be obtained by replacing y_i by x_{ij} in ϕ_0 and y'_j by x_{ij} in ψ_0 . Thus \hat{y} and \hat{y}' are \bar{y} and \bar{y}' with y_i or y'_j deleted. Write \hat{x} for \bar{x}, x_{ij} .

Let $\phi^* = \exists \hat{y} \phi_0^*(\hat{x}, \hat{y})$ and $\psi^* = \exists \hat{y}' \psi_0^*(\hat{x}, \hat{y}')$. Then ϕ^* has n quantified variables and $T_C \models \neg(\phi^* \wedge \psi^*)$, since any model of $T_C \cup \{\phi^*, \psi^*\}$ gives rise to a model of $T_C \cup \{\phi, \psi\}$; the variables y_i, y'_j may be realized by the value of x_{ij} .

By induction hypothesis for each choice of i and j , ψ^* contains a subformula ψ_{ij}^* involving at most b_n variables so that $T_C \models \neg(\phi^* \wedge \psi_{ij}^*)$.

Let $\bar{y}'' \subseteq \bar{y}'$ be the set of at most $k + k^2 b_n$ variables consisting of $C \cap \bar{y}'$ together with the all quantified variables from any ψ_{ij}^* , and let ψ_1 be the restriction of ψ to \bar{y}'' . Then we claim

$$T_C \models \neg(\phi \wedge \psi_1) \quad (*)$$

so we may take $b_{n+1} = k + k^2 b_n$.

For (*), consider any model \mathcal{M} of $\phi \wedge \psi_1$. Then C embeds in the free join over \bar{x} of the induced graphs A, B on \bar{x}, \bar{y} and \bar{x}, \bar{y}'' . So if \mathcal{M} omits \mathcal{C} , there must be some identification $x_i = y_j$ with $x_i, y_j \in C$. This is exactly what is ruled out by ψ_{ij}^* . ■

Corollary to Theorem 1 *Let \mathcal{C} be a finite class of finite graphs. Then T_C^* is model complete and is the model companion of T_C .*

Proof:

This is equivalent to Theorem 1, part (1). ■

§2. Universal Graphs and existential types.

In this section we give criteria for the existence of a universal graph in \mathcal{G}_C , for C a finite set of finite connected graphs. We will show that when there is a universal graph in \mathcal{G}_C , there is a *canonical* one, namely the “ \aleph_0 -saturated” graph in \mathcal{E}_C . We will also show the relationship of this problem to a model theoretic notion of *algebraic closure*. We review the definitions.

Definition 3

Let \mathcal{C} be a collection of finite forbidden subgraphs.

1. The *existential type* $tp_G(\bar{a})$ of a finite sequence $\bar{a} = a_1, a_2, \dots, a_n$ in a graph $G \in \mathcal{E}_C$ is the set of existential formulas $\phi(\bar{x})$ such that $G \models \phi(\bar{a})$. The Stone space $S_n(T_C^*)$ is the set of all existential types $tp(\bar{a})$ of sequences $\bar{a} = a_1, \dots, a_n$ in any graph $G \in \mathcal{E}_C$.

2. $G \in \mathcal{E}_C$ is \aleph_0 -saturated if for all n , all $\bar{a} \in G$ of length n , and all $(n + 1)$ -types in $S_{n+1}(T_C^*)$ whose restriction to the first n variables is $tp_G(\bar{a})$, there is $v \in V(G)$ so that $tp_G(\bar{a}, v)$ is the specified type.

Example 7 When $\mathcal{C} = \{S_3\}$, specifying the type of an element a in $G \in \mathcal{E}_{\mathcal{C}}$ is equivalent to describing the isomorphism type of its connected component in G . In particular if a_1, \dots, a_n lie in distinct components isomorphic to 2-way infinite paths, ω -saturation yields an element a_{n+1} lying in another component isomorphic to such a path. Thus the \aleph_0 -saturated model is the largest model in $\mathcal{E}_{\mathcal{C}}$. This is the case in general.

Theorem 2 Let \mathcal{C} be a finite set of connected forbidden subgraphs. Then the following are equivalent:

- 1). There is a universal graph in $\mathcal{G}_{\mathcal{C}}$.
- 2). There is a universal graph in $\mathcal{E}_{\mathcal{C}}$.
- 3). $\mathcal{E}_{\mathcal{C}}$ contains a unique \aleph_0 -saturated graph, up to isomorphism.
- 4). $S_n(T_{\mathcal{C}}^*)$ is countable, for any n .

Proof:

This is a special case of general model theoretic principles [CK, §2.3]. We sketch the ideas here.

The equivalence of 1) and 2) is immediate. It suffices to note that any $G \in \mathcal{G}_{\mathcal{C}}$ embeds into a $G^* \in \mathcal{E}_{\mathcal{C}}$. For the equivalence of 2) to 4) one recalls that $\mathcal{E}_{\mathcal{C}}$ is the class of countable models of $T_{\mathcal{C}}^*$. We will show $2) \Rightarrow 4) \Rightarrow 3) \Rightarrow 2)$.

$2) \Rightarrow 4)$. Let $G \in \mathcal{E}_{\mathcal{C}}$ be universal. As G is countable, the set $\{tp_G(\bar{a}) : \bar{a} = a_1, \dots, a_n \in G\}$ is countable. Any type $tp_{G'}(\bar{a})$ realized in any $G' \in \mathcal{E}_{\mathcal{C}}$ will be realized in G since we may take G' to be an induced subgraph of G by universality and $tp_G(\bar{a}) = tp_{G'}(\bar{a})$ by existential completeness.

$4) \Rightarrow 3)$. If $S_n(T_{\mathcal{C}}^*)$ is countable for all n , one builds a countable saturated model as the limit of an increasing countable sequence of models in $\mathcal{E}_{\mathcal{C}}$, see [CK, Theorem 2.3.7].

The uniqueness follows from the completeness of $T_{\mathcal{C}}^*$ [CK, Theorem 2.3.7].

$3) \Rightarrow 2)$. Saturated models are universal [CK, Theorem 2.3.10]. ■

In the examples, one often encounters the special case in which $\mathcal{E}_{\mathcal{C}}$ contains a unique model up to isomorphism, so that the \aleph_0 -saturation condition is vacuous. This is a rather special case in model theory, and the frequency of its occurrence in our context is an indication that something more specialized is involved. To analyze this further we introduce the notion of *algebraic closure*.

Definition 4

Let \mathcal{C} be a set of forbidden subgraphs, $G \in \mathcal{E}_{\mathcal{C}}$, $A \subseteq G$, $a \in G$. We say that a is algebraic over A (in G) if there is an existential formula $\phi(x, \bar{a})$ with $\bar{a} \in A$ such that the set $\{a' \in G : \phi(a', \bar{a})\}$ is finite and contains a . We write $\text{acl}_G(A)$ (algebraic closure) for the set of $a \in G$ algebraic over A . We say A is algebraically closed in G if $\text{acl}_G(A) = A$.

Lemma 3 Let \mathcal{C} be a finite set of connected forbidden subgraphs. If $\mathcal{G}_{\mathcal{C}}$ contains a universal graph then the set of isomorphism types of induced subgraphs of graph $G \in \mathcal{E}_{\mathcal{C}}$ on subsets of the form $\text{acl}(A)$ with A finite, is countable.

Proof:

Let $G \in \mathcal{E}_{\mathcal{C}}$ be universal. Then for any $G' \in \mathcal{E}_{\mathcal{C}}$ and any $A \subseteq G'$ finite, an embedding ι of G' into G given an isomorphism between $G' \upharpoonright \text{acl}_{G'}(A)$ and $G \upharpoonright \text{acl}_G(\iota A)$. The point here is that $\text{acl}_G(\iota A) = \iota[\text{acl}_{G'}(A)]$, by existential completeness. ■

It would be pleasant if the converse held: in other words, to show the nonexistence of universal graphs one would be obligated to construct uncountably many isomorphism types of algebraic closures of finite sets. This is what has actually occurred in all examples treated to date [ChK,CS2,CST,FK1,Ko,FK2,GK,KP1].

In fact, in most cases one of the two extremes of the following pseudo-dichotomy have been encountered:

I. The algebraic closure of a finite set is finite.

II. There are uncountably many isomorphism types of induced subgraphs on sets $\text{acl}_G(A)$, with A finite, in graphs $G \in \mathcal{E}_C$.

On the other hand the example $\mathcal{C} = \{S_3\}$, a star of degree 3, shows that case I is indeed a special case, as one might anticipate. This makes it all the more surprising that this case is typical in practice, in contexts where universal graphs exist.

All of this leaves open the possibility, already referred to, that case II is an exact criterion for the nonexistence of universal graphs. To refute this in the category of graphs is not so easy. We will give an example in the category of vertex-colored graphs. It should not be too difficult to encode this as an example in the category of graphs, but it would be more to the point to prove the general encoding conjecture noted in the introduction, which we will not undertake here.

Example 8 *We work with vertex colored graphs in which there are three colors: 0, +1, -1. Each vertex of color 0 has at most two neighbors of color 0, and only one of the other two colors occurs among its neighbors. Vertices of colors +1 and -1 are adjacent to at most one vertex, which must have color 0. This clearly corresponds to a finite set of connected forbidden subgraphs. In \mathcal{E}_C the graphs consist of cycles and 2-way infinite paths made up of vertices of color 0, each decorated with infinitely many adjacent vertices of color +1 or -1. It is easy to see that the algebraic closure of a finite set A consists of the union of the connected components of vertices of color 0 in or adjacent to A , together with vertices in A of color +1 and -1. Thus there are countably many induced subgraphs on $\text{acl}(A)$ for A finite.*

However, the type of an element v of color 0 contains a specification of the colors +1 and -1 of the neighbors of all vertices of color 0 in its connected component. Thus $S_1(T_C^)$ is uncountable. It follows that the types in general contain information not controlled by the algebraic closure operation.*

On the other hand, we will show that when condition (I) holds, control of algebraic closure is enough. Indeed, in the example just discussed, there are only countably many types associated with vertices of color 0 whose connected component, among the vertices of type 0, is finite. In fact, if the order of the connected component in question is specified, there are finitely many possible types.

§3. \aleph_0 -categoricity and local finiteness

A theory is said to be \aleph_0 -categorical if it has a unique countable model, up to isomorphism. As we have noted, among theories of the form T_C^* for which a universal countable model exists, the \aleph_0 -categorical case is surprisingly common. The next result casts some light on this phenomenon.

Theorem 3 *Let \mathcal{C} be a finite set of connected finite graphs. Then the following are equivalent:*

- (1). T_C^* is \aleph_0 -categorical.
- (2). $S_n(T_C^*)$ is finite for each n .
- (3). For $A \subseteq \mathcal{M} \models T_C^*$ finite, we have $\text{acl}(A)$ finite.

These conditions imply

(4). \mathcal{G}_C contains a universal countable graph.

By Theorem 1, (2), T_C^* is complete. Therefore the equivalence of (1) and (2) holds by general model theory [CK, Theorem 2.3.13.]. That (1) implies (4), and (2) implies (3), are both immediate. Thus all that requires proof is the implication from (3) to (2). For this we prove a more refined technical lemma, based on the following definition and fact.

Definition 5 *Let G be a graph, and $A \subseteq G$. Set*

$$\begin{aligned} \text{tp}_n(A) = \{ \phi(\bar{a}) : \phi \text{ is existential, with at most } n \text{ quantified variables,} \\ \bar{a} \in A, \text{ and } \phi(\bar{a}) \text{ holds in } G \}. \end{aligned}$$

(This depends on G , and one may write $\text{tp}_n^G(A)$ to show this dependence.)

Fact 1. (Park, cited in [Ba]) *Let A be algebraically closed in B . Then there is $C \succ B$ and $B' \simeq B$ (over A) with $B' \prec C$ and $A = B \cap B'$. Note that in [Ba] the term “Park-a.c.” is used for our “algebraically closed”.*

Lemma 4 *Let \mathcal{C} be a finite set of finite graphs, and $A \subseteq G \in \mathcal{E}_C$ with A finite and algebraically closed. Then for $n = \max\{|C| : C \in \mathcal{C}\}$, $\text{tp}_n^G(A)$ determines $\text{tp}(A)$.*

Proof:

We write \bar{a} for A arranged as a finite sequence. Let $e(\bar{a})$ be an existential sentence. We claim that $e(\bar{a})$ holds in G if and only if the following theory T_e is consistent:

$$(T_e) \quad \text{“} A \text{ is algebraically closed”} \cup T_C \cup \text{tp}_n^G(\bar{a}) \cup \{e(\bar{a})\}.$$

One may easily find axioms expressing the assertion that A is algebraically closed. Thus T_e is indeed a first order theory. If $e(\bar{a})$ holds in G , then T_e holds in G and thus T_C is consistent.

Suppose conversely that T_e holds in some G_1 . We claim that $e(\bar{a})$ will then hold in G .

Let $e(\bar{a}) = \exists \bar{y} e_0(\bar{a}, \bar{y})$ with e_0 quantifier-free. We may suppose that e is primitive, and e_0 is conjunctive. Choose \bar{b} in G_1 so that $e_0(\bar{a}, \bar{b})$ holds. We may suppose $\bar{b} \cap A = \emptyset$, adjusting e_0 if necessary. Form $G' = G \cup \bar{b}'$ by freely amalgamating G with a copy $\bar{a}\bar{b}'$ of $\bar{a}\bar{b}$ over \bar{a} . That is, the edges in G' lie in G and in $\bar{a}\bar{b}'$. Note that G and G_1 agree on \bar{a} , as a description of the induced graph on \bar{a} is contained in $\text{tp}_n(\bar{a})$.

If $G' \in \mathcal{G}_C$ then as $G \subseteq G'$, $G \in \mathcal{E}_C$, and $e(\bar{a})$ holds in G' we find that $e(\bar{a})$ holds in G , as claimed. Suppose now that $G' \notin \mathcal{G}_C$. Then we will show that $G_1 \notin \mathcal{G}_C$, contradicting our assumption on G_1 .

We have some $C \in \mathcal{C}$ which embeds into G' , and we may take $C \subseteq G'$. Let $\bar{c}_0' = C - (A \cup \bar{b}') \subseteq G$, and let $\phi_0(\bar{a}, \bar{c}_0')$ be a conjunctive quantifier-free formula specifying the isomorphism type of the induced subgraph on \bar{a}, \bar{c}_0' . Then the existential formula

$$\phi(\bar{x}) = \text{“}\exists \bar{y} \phi_0(\bar{a}, \bar{y})\text{”}$$

belongs to $\text{tp}_n^G(A)$. Hence we have \bar{c}_0 in G_1 satisfying $\phi_0(\bar{a}, \bar{c}_0)$.

As $\bar{c}_0 \cap A = \emptyset$ and A is algebraically closed in G_1 , by repeated applications of Fact 1 we can find disjoint sequences $\bar{c}_0^{(1)}, \dots, \bar{c}_0^{(k)}$ in G_1 , for any k , so that the induced subgraphs on $\bar{a}\bar{c}_0^{(i)}$ are isomorphic to $\bar{a}\bar{c}_0$ in the natural order.

Choose $k > |\bar{b}|$. Then for some i , $\bar{c}_0^{(i)} \cap \bar{b} = \emptyset$, and thus the free amalgam of $\bar{a}\bar{c}_0$ with $\bar{a}\bar{b}$ over \bar{a} embeds into the induced graph on $\bar{a}\bar{b}\bar{c}_0^{(i)}$. But this free amalgam is also isomorphic to the subgraph of G' induced on $\bar{a}\bar{b}'\bar{c}_0'$, which is C . Thus C embeds in G_1 , a contradiction. ■

Proof of Theorem 3:

As noted above, we need only check $(3) \Rightarrow (2)$. Assuming (3), then for n fixed there is a uniform bound on $|A|$ for A the algebraic closure of a set of n elements in a model of $T_{\mathcal{C}}^*$. Thus it suffices to show that for each such A , the type of A in a graph $G \in \mathcal{E}_{\mathcal{C}}$ is determined up to finitely many possibilities. Indeed, with A fixed, by the preceding lemma there is N such that $\text{tp}_N(A)$ determines $\text{tp}(A)$; and there are only finitely many possibilities for $\text{tp}_N(A)$. ■

Thus if the algebraic closure operation is uniformly locally finite on $\mathcal{E}_{\mathcal{C}}$, a universal graph exists. Earlier we showed by example that when it is not uniformly locally finite, knowledge of this operator does not in general settle the question of the existence of a universal graph: of course, at the other extreme (case II of §2), the question is also settled by the structure of algebraic closure.

§4. Algebraic closure.

In view of the importance of the algebraic closure operator in dealing with problems of universality, it is worth while making explicit what is involved.

Definition 6 *Let A, B be graphs and $f : V(A) \rightarrow V(B)$. Then f is a homomorphism if f carries edges to edges.*

Remarks

1. An injective homomorphism is an isomorphism with a subgraph (not necessarily induced).
2. We deal throughout with graphs without loops. In particular if a homomorphism $f : A \rightarrow B$ identifies two vertices of A , they cannot be linked by an edge. (We could just as well allow loops. In this case, if the loop on one vertex is in \mathcal{C} , we recover the loop-free context.)

Lemma 5 *Let \mathcal{C} be a finite collection of finite graphs, and $A \subseteq G \in \mathcal{E}_{\mathcal{C}}$. Then the following are equivalent:*

- (1). A is not algebraically closed in G .
- (2). There is some $C \in \mathcal{C}$ and a homomorphism $C \rightarrow C' \subseteq G$ so that C embeds in the free amalgam over A of $|C|$ copies of C' .

Proof:

(2) \Rightarrow (1) : Let $h : C \rightarrow C'$ as in (2) and let $B = C' - A$. If $G - A$ contains $|C|$ disjoint copies B^i of B (isomorphic over A), then the free amalgam of $|C|$ copies of C' over A embeds into $A \cup \bigcup_{i \leq |C|} B^i$, and hence C embeds in G , a contradiction. By Park's Theorem (Fact 1), our claim follows.

(1) \Rightarrow (2): As A is not algebraically closed, there is $b \in \text{acl}(A) - A$, and there is an existential formula $\phi(\bar{a}, b) = \text{"}\exists \bar{y} \phi_0(\bar{a}, b, \bar{y})\text{"}$ so that $|\{b' \in G : \phi(\bar{a}, b')\}| = k < \infty$. Let $\bar{b} \in G$ satisfying $\phi_0(\bar{a}, b, \bar{b})$, and set $B = \{b\} \cup \bar{b}$. With a slight change of notation, we may suppose $B \cap A = \emptyset$.

Let $G_0 = AB^1 \cdots B^{k+1}$ be the free amalgam over A of $k+1$ copies AB^i of AB (isomorphic over A). Let G_1 be the free amalgam of G and G_0 over A . Then $G_1 \notin \mathcal{G}_{\mathcal{C}}$, as otherwise after extending to $G_2 \in \mathcal{G}_{\mathcal{C}}$, we find $G \prec G_2$ but $|\{b' \in G_2 : \phi(\bar{a}, b')\}| > k$, a contradiction.

As $G_1 \notin \mathcal{G}_{\mathcal{C}}$, there is $C \in \mathcal{C}$ and an embedding $f : C \rightarrow G_1$. We alter this to a homomorphism $h : C \rightarrow G$ by mapping each B^i isomorphically over A to B . Let C' be the image of h . Then the free join

of $|C|$ copies of $C' - A$ over A contains the image of f , as required. \blacksquare

We give a simple example to illustrate the power of this result. Later as we go into applications in more detail, we will get considerably more mileage out of the same idea.

Theorem 4 *Let \mathcal{C} be a finite set of connected finite graphs. Suppose that for any $C \in \mathcal{C}$ and any surjective homomorphism $h : C \rightarrow C'$, that C' contains a graph in \mathcal{C} . Then for $A \subseteq G \in \mathcal{E}_{\mathcal{C}}$, $\text{acl}(A) = A$. In particular, $T_{\mathcal{C}}^*$ is \aleph_0 -categorical and hence there is a universal graph in $\mathcal{G}_{\mathcal{C}}$.*

Proof:

If $A \subseteq G$ is not algebraically closed, application of Lemma 5 produces $h : C \rightarrow C' \subseteq G$, but then $G \notin \mathcal{G}_{\mathcal{C}}$. \blacksquare

Example 9 *Fix k . Let \mathcal{C} consist of all cycles of odd lengths, up to $2k + 1$. Then there is a universal graph in $\mathcal{G}_{\mathcal{C}}$.*

This result was first proved in [KMP] with an elaborate amalgamation argument, containing some minor inaccuracies which were subsequently corrected. This should serve to illustrate the utility of our general considerations. We will use the same idea below to construct a number of new examples.

II. Applications

In this part, we first review the known results from our point of view. From this point of view, the main question is the behavior of the algebraic closure operation on finite sets. This qualitative problem can be rephrased in quantitative terms; from that point of view, the known results leave open a number of questions regarding the estimates for the size of $\text{acl}(A)$ in terms of $|A|$, and similar issues, which we will point out in detail.

§5. Negative results: explosion of algebraic closure.

The negative results all depend on the construction of 2^{\aleph_0} nonisomorphic induced graphs of the form $\text{acl}(A)$ for A of some fixed size, which can be read off explicitly from the various papers, though the terminology varies somewhat. In such cases there are two natural questions concerning $|\text{acl}(A)|$:

(I). What is the least cardinality α such that there are 2^{\aleph_0} possible isomorphism types for the graph induced on $\text{acl}(A)$ in a graph $G \in \mathcal{E}_{\mathcal{C}}$, with $|A| = \alpha$?

(II). What is the least cardinality α' such that $\text{acl}(A)$ is infinite in some $G \in \mathcal{E}_{\mathcal{C}}$, with $|A| = \alpha'$?

One suspects these are usually equal, though exceptions were mentioned earlier. All of the negative results on universal graphs to date may be phrased as explicit upper bounds on α in various cases.

5.1.2-connected graphs.

The main result of [FK1] gives a bound for α when $\mathcal{C} = \{C\}$ consists of a single constraint C which is 2-connected and not complete:

$$(1) \quad \alpha \leq 4(2k)^{2g}.$$

with $k = 2|V(G)| - 1, g = |V(G)| + 1$. Actually the result proved is significantly more general. The same bound is given when C contains a block C_0 which is 2-connected, and which contains two nonadjacent vertices u, v so that C_0 does not embed in C_{uv} , the graph obtained from C by identifying u with v (keeping all edges).

Various special cases proved earlier give sharper estimates for more specific constraints. In [KP1] one finds $\alpha \leq 8m - 7$ when $C = K_{m,n}$ is complete bipartite ($m \leq n$). In [ChK] one finds $\alpha \leq 4N + 1$ with $N = (14^\nu - 1)/13$, $\nu = |V(G)|$, for G a cycle of length at least 4; and the same bound is obtained in [CS2] when \mathcal{C} is a finite set of cycles, taking $\nu = \max\{|V(G)| : C \in \mathcal{C}\}$. There is one exception in this case: when \mathcal{C} consists of all odd cycles up to some bound, there is a universal graph ($\alpha = \infty$); this was mentioned above, following Theorem 4.

The special case in which \mathcal{C} consists of all cycles up to some even bound was considered in [GK]; they found $\alpha \leq 5$ in this case.

All of this raises a number of natural questions. First of all, can one combine [FK1] and [CS2] to identify all finite sets \mathcal{C} of 2-connected graphs for which there is a corresponding universal graph, and to estimate α in the other cases?

Secondly, can one obtain a respectable lower bound for α , or at least determine whether α is unbounded in most cases? Some information is provided by the following:

Lemma 6 *Let \mathcal{C} be a finite set of k -connected graphs. Let $G \in \mathcal{E}_{\mathcal{C}}$ and $A \subseteq V(G)$, $|A| < k$. Then $\text{acl}(A) = A$. In particular $\alpha \geq k$.*

Proof:

We apply Lemma 5. If $C \in \mathcal{C}$ embeds in a free amalgam of copies of C' over A , then as C is k -connected with $k > |A|$, C would embed in C' , hence in G . ■

Example 10 1. If \mathcal{C} is a finite set of cycles, this tells us only that $\text{acl}(A) = A$ when $|A| = 1$.

2. If $\mathcal{C} = \{K_{m,n}\}$ with $m \leq n$ we find $\text{acl}(A) = A$ for $|A| < m$, and $\alpha \geq m$. This matches the upper bound in [KP1] reasonably well.

3. If $\mathcal{C} = \{C\}$ with C a complete graph K_n with one edge deleted, we find $\text{acl}(A) = A$ for $|A| \leq n - 2$, so $\alpha \geq n - 1$.

This leaves a rather large gap between the upper and lower bounds for α in most cases. One suspects the upper bounds could be sharpened considerably.

5.2. Trees.

A tree is called *bushy* if it has no vertices of degree 2. For constraint sets \mathcal{C} consisting of a single bushy tree with at least 5 vertices, the result of [CST] yields a bound slightly sharper than the following:

$$(2) \quad \alpha < n.$$

This can be radically improved: if $\alpha < \infty$ then $\alpha = 1$ in the case of trees (Proposition 6 below).

One peculiarity of tree constraints is that for any $G \in \mathcal{E}_{\mathcal{C}}$ (where \mathcal{C} consists of a single tree constraint T) we never have $\text{acl}(A) = A$ when $|A| = 1$, unless $|V(T)| = 2$. This can be seen using Lemma 5.

§6. Positive results: local finiteness estimates

In most of the cases in which a universal graph is known to exist, $T_{\mathcal{C}}^*$ is \aleph_0 -categorical, and the situation is described by Theorem 3. In such cases the criterion in part (3) of that theorem has not been used. Indeed there are a variety of approaches to \aleph_0 -categoricity and it does not seem reasonable to insist on one as most appropriate in all cases, but we have indicated some situations in which the computation of algebraic closure is effective, following Theorem 4. We will give some new applications in the following section. Here we

review the known positive results, with an eye on the additional information they furnish about algebraic closure in such cases. The natural problem here is to estimate the function

$$c(n) = \max\{|\text{acl}_G(A)| : |A| = n, A \subseteq V(G), G \in \mathcal{E}_C\}.$$

Here upper bounds are the main point, but one may look for accurate asymptotics.

6.1. Trees.

Tallgren has conjectured that the only trees T for which $\mathcal{G}_{\{T\}}$ has a universal object are the paths and the trees obtained from a path by attaching one additional edge. His proof of the existence of a universal graph in the latter case is unpublished, but this case is of considerable interest as it affords an example in which T_C^* is not \aleph_0 -categorical, but a universal graph exists. This point is illustrated quite well by the simple example of a star S_3 of degree 3, discussed as Example 4 in §1.

6.2.2-connected case.

Previously only two examples of finite sets \mathcal{C} of 2-connected graphs were known for which \mathcal{G}_C has a universal graph: $\mathcal{C} = \{K_n\}$, a single complete graph, or $\mathcal{C} = \{C_{2k+1} : k \leq n\}$ the set of odd cycles of size up to to some bound. Both are covered by Theorem 4, as noted earlier: indeed $\text{acl}(A) = A$ for all A , and $c(n) = n$. Additional examples arising from Theorem 4 will be considered in the next section.

6.3. Bow-Ties

Any graph can be analyzed as constructed from a tree of “blocks” (2-connected graphs and edges). However, we know of no way to combine the analysis of 2-connected constraints and tree constraints to produce something more general. For that matter, relatively few explicit examples have been successfully analyzed to date. Komjáth [Ko] did find one example in which a universal graph exists. Such examples are presumably quite rare. We will give new examples in §8. Here we give an analysis of the Füredi-Komjáth “bow-tie” example in terms of our machinery of algebraic closure.

A *bow-tie* is the graph on five vertices formed by attaching two triangles to a common vertex. More generally, one may consider bouquets of complete graphs with one common vertex. For bouquets of at least three complete graphs of constant size, it is shown in [Ko] that the only ones corresponding to universal graphs are the bow-tie and the degenerate bouquets consisting of one complete graph.

Let B be the bow-tie, $\mathcal{C} = \{B\}$. We show that T_C^* is \aleph_0 -categorical, and in particular there is a universal countable bow-tie-free graph. This follows by combining Theorem 3 with the following estimate.

Proposition 1 *Let $G \in \mathcal{E}_{\{B\}}$, $A \subseteq G$ finite. Then $|\text{acl}(A)| \leq 4|A|$.*

Proof:

Call an edge of G *special* if it lies in two triangles of G . We make the following claims, which will be verified below.

- (1). Every triangle in G contains at least one special edge.
- (2). Every point that lies on a triangle, but no special edge of that triangle, lies on a unique triangle.
- (3). If a point lies on two special edges, it lies on a graph $K \cong K_4$. In this case, any triangle containing that point is contained in K .

Assuming these claims for the moment, we proceed as follows. Given $G \in \mathcal{E}_{\{B\}}$, $A \subseteq V(G)$ finite, let A^* be the union of A with the set of all vertices of G which lie on special edges which themselves lie on triangles

containing a point of A . It follows from (2),(3) that

$$(4) \quad |A^*| \leq 4|A|.$$

Thus it will suffice to show that A^* is algebraically closed. We show first

$$(5) \quad A^{**} = A^*.$$

Let $u \in A^{**} - A$. Then u lies on a special edge e , where e lies on a triangle t meeting A^* . Let the vertices of e be $\{u, v\}$, and let the third vertex of t be w .

We claim $u \in A^*$. Assume not. Then v or w belongs to A^* , but t contains no vertex of A . Thus v or w is in $A^* - A$.

If $v \in A^* - A$, then v lies on a special edge e' which lies on a triangle t' meeting A in a vertex a . If $e' = e$ this forces $u \in A^*$, as desired. If $e' \neq e$ then by (3) v lies on a K_4 , containing t and t' . Hence a, u, v are the vertices of a triangle in G and therefore $u \in A^*$, as claimed.

If $w \in A^* - A$ and w lies on a special edge of t , then by (2) all edges of t are special, and then the argument above applies to w . If not, then by (2) w lies on a unique triangle. Then if $w \in A^*$ then t meets A , a contradiction.

Now we show

$$(6) \quad A^* \text{ is algebraically closed.}$$

We apply Lemma 5. B has only two proper homomorphic images, so applying the criterion of Lemma 5, if A^* is not algebraically closed then there is a triangle t meeting A^* in one vertex. By (1) t contains a special edge, and some vertex of that edge then lies in $A^{**} - A^*$, a contradiction.

It remains to verify our claims (1)-(3). Both (2) and (3) are direct consequences of the assumption that G is B -free, by inspection. We turn to (1).

Let e be an edge of a triangle t lying in G . Let G^* be the graph formed from G by attaching an additional triangle containing the edge e . In G^* , e is special. If G^* is B -free, then as $G \in \mathcal{E}_{\{B\}}$, e is special in G . If G^* is not B -free, then as G is B -free, it follows that one of the other two edges of t is special in G . ■

§7. New universal graphs.

We gave a general construction in Theorem 4 which produces finite sets \mathcal{C} of connected constraints for which $T_{\mathcal{C}}^*$ is \aleph_0 -categorical and hence, in particular, there is a universal \mathcal{C} -free graph. We now generalize this.

Theorem 5 *Let \mathcal{C} be a finite set of finite connected graphs such that $T_{\mathcal{C}}^*$ is \aleph_0 -categorical. Let \mathcal{H} be a finite set of finite connected graphs which is closed under homomorphic image. Then $T_{\mathcal{C} \cup \mathcal{H}}^*$ is \aleph_0 -categorical.*

Proof:

Let $G \in \mathcal{E}_{\mathcal{C} \cup \mathcal{H}}$, and $A \subseteq G$ finite. We must bound $|\text{acl}(A)|$ in terms of $|A|$ and apply Theorem 3. Let $G' \in \mathcal{E}_{\mathcal{C}}, G \subseteq G'$. Let $B = \text{acl}_{G'}(A)$, a finite set of size bounded by a function of $|A|$. It suffices to show that $B_0 = B \cap G$ is algebraically closed in G .

Suppose the contrary, by Lemma 5 we have some $C \in \mathcal{C} \cup \mathcal{H}$ and a homomorphic image $C' \subseteq G$ so that C embeds in the free amalgam over B_0 of $|C|$ copies of C' . If $C \in \mathcal{H}$ then $C' \in \mathcal{H}$, contradicting the assumption that G omits $\mathcal{C} \cup \mathcal{H}$. Hence $C \in \mathcal{C}$.

Now we consider B_0, B , and C' in G' . As G' omits C , G' does not contain the free amalgam of $|C|$ copies of C' over B_0 or over B ($C' \cap B = C' \cap B_0$). Then by Park's theorem, B is not algebraically closed in G' , a contradiction. ■

It is of course trivial to produce examples of set of constraints to play the role of \mathcal{H} here, but we will want to consider a number of concrete constructions, particularly with a view toward keeping $|\mathcal{H}|$ small. This will require some preliminary observations.

Remarks

1. We do not in fact require \mathcal{H} to be *closed* under the formation of homomorphic images. What is needed is the following: if C' is a homomorphic image of $C \in \mathcal{H}$, then C' contains an element of \mathcal{H} . In the future we will take this condition as the *definition* of “closure under homomorphism”.

2. In particular if A is a finite connected graph we will write $\mathcal{H}_0(A)$ for the set of all homomorphic images of A and $\mathcal{H}(A)$ for the set of minimal elements of $\mathcal{H}_0(A)$ (with respect to embeddings as subgraphs). For example, if A is a cycle of odd length $2N + 1$, then $\mathcal{H}(A)$ consists of odd cycles of length $2n + 1, n \leq N$. Similarly, if A is a bipartite graph containing at least one edge, then $\mathcal{H}(A) = \{K_2\}$. More general, for any finite connected graph A , $\mathcal{H}(A)$ contains a unique complete graph K_n , with $n = \chi(A)$ the chromatic number. Thus one only gets new examples by considering graphs of chromatic number χ which do not contain the complete graphs K_χ . In this case $|\mathcal{H}(A)| \geq 2$.

Definition 7 1. Let A_1, A_2 be two graphs. Then $A_1 \times A_2$ is the graph with vertex set $V(A_1) \cup V(A_2)$, and whose edges are those of A_1 and A_2 together with all pairs (u, v) , where $u \in A_1, v \in A_2$ or vice versa.

2. Let $\mathcal{C}_1, \mathcal{C}_2$ be two sets of graphs. Then $\mathcal{C}_1 \times \mathcal{C}_2 = \{A_1 \times A_2 : A_1 \in \mathcal{C}_1, A_2 \in \mathcal{C}_2\}$.

Remark

$$\mathcal{H}_0(A_1 \times A_2) = \mathcal{H}_0(A_1) \times \mathcal{H}_0(A_2).$$

Example 11 With M, N fixed integers, the class of graphs omitting $C_{2m+1} \times C_{2n+1}$ for $m \leq M, n \leq N$ has a universal graph; in particular for $M = 0$, this is the class constrained by forbidding “wheels” $\{K_1 \times C_{2n+1} : n \leq N\}$.

Another family of well-behaved examples is generated by application of a construction used by Mycielski to generate triangle free graphs of arbitrary high chromatic number, where Mycielski would begin with K_2 , we substitute K_n , getting the following graphs, which we call M_n . Let $V(M_n) = \{0\} \cup (\{1, 2, \dots, n\} \times \{0, 1\})$ and set $A_n = \{1, 2, \dots, n\} \times \{0\}, B_n = \{1, 2, \dots, n\} \times \{1\}$. Edges are defined as follows. The vertex 0 is adjacent to the vertices of A_n and no others; M_n induces a complete graph on B_n , and no edges on A_n ; and the vertices $(i, 0)$ and $(j, 1)$ are adjacent if and only if $i \neq j$. This graph arises by applying Mycielski's construction to B_n (i.e. K_n). To have a more suggestive notation we write a_i for $(i, 0)$ and b_i for $(i, 1)$.

Lemma 7 $\mathcal{H}(M_n) = \{K_k \times M_{n-k} : k \leq n, k \neq n - 1\}$. In particular $\chi(M_n) = n + 1$ and $|\mathcal{H}(M_n)| = n$.

Proof:

Let $\mathcal{C} = \{K_k \times M_{n-k} : k \leq n\}$. Then $\mathcal{C} \subseteq \mathcal{H}_0(M_n)$. To see this, identify the vertices a_i and b_i for $i > n - k$. In particular for $k = n$ we have $K_n \times M_0 = K_{n+1} \in \mathcal{C}$, so $\chi(M_n) \leq n + 1$, and $\chi(M_n) \geq n + 1$ by

Mycielski's argument [BM, §8.5].

Any homomorphic image of M_n other than those listed will involve either the identification of the vertex 0 with a vertex of B_n , or the identification of vertices in A_n . In either case the resulting homomorphic image contains K_{n+1} by inspection. Thus the minimal homomorphic images of M_n belong to $\mathcal{C} : \mathcal{H}(M_n) \subseteq \mathcal{C}$. Furthermore for $k = n - 1, K_k \times M_{n-k} = K_{n-1} \times M_1 \supseteq K_{n-1} \times K_2 = K_{n+1}$, so $K_{n-1} \times M_1 \notin \mathcal{H}(M_n)$.

It remains to be shown that the graphs $K_k \times M_{n-k}$ for $k \leq n, k \neq n - 1$ are incomparable; this will complete the characterization of $\mathcal{H}(M_n)$.

Suppose therefore that $K_k \times M_{n-k}$ embeds in $K_l \times M_{n-l}$ with $0 \leq k, l \leq n$ and $k, l \neq n - 1, k \neq l$. As $|K_k \times M_{n-k}| \leq |K_l \times M_{n-l}|$ we have $k \geq l$. The case $k = n$ may be eliminated by inspection. Accordingly we assume $0 \leq l < k \leq n - 1$.

Let $f : K_k \times M_{n-k} \rightarrow K_l \times M_{n-l}$ be an embedding. As $k > l$, fix $u \in K_k$ so that $f(u) \notin K_l$. Now u has $2n - k$ neighbors in $K_k \times M_{n-k}$. If $f(u) \notin B_{n-l}$, then $f(u)$ has at most n neighbors in $K_l \times M_{n-l}$, forcing $2n - k \leq n$, a contradiction. So $f(u) \in B_{n-l}$ and as u is adjacent to every other vertex of $K_k \times M_{n-k}$, $f[K_k \times M_{n-k}]$ does not contain the vertex labelled 0 in $K_l \times M_{n-l}$. However the graph resulting from deletion of this vertex has chromatic number n , while $K_k \times M_{n-k}$ has chromatic number $n + 1$, a contradiction. ■

Examples of constraint sets \mathcal{C} allowing a universal graph with $|\mathcal{C}| = 1$ are very rare, and indeed few examples are known with any sharp bound on $|\mathcal{C}|$. We will consider the possibilities in the case $\mathcal{C} = \mathcal{H}(A)$. Evidently, if we require $|\mathcal{C}| = 1$ we will have $\mathcal{C} = \{K_n\}$ for some complete graph, which is one of the oldest examples. We can on the other hand produce a number of new examples with $|\mathcal{H}(A)| = 2$. We note first the simple example $\mathcal{H}(C_s \times K_n) = \{C_s \times K_n, K_{n+3}\}$. It seems possible *a priori* that these are the only such examples, and we therefore will give an additional construction, showing at least that it will not be easy to classify the cases with $|\mathcal{H}(A)| = 2$.

Construction Let $G = r \cdot K_n + K_m$ be the disjoint sum of r complete graphs K_n , and one more, K_m , with $n \geq m \geq 1$ and either $r \geq 2$, or $m \geq n - 1$.

We will write $G = A_1 + \dots + A_r + B$ with $A_i \simeq K_n$, $B \simeq K_m$. Let $m_0 = \min(m, n - 1)$ and let $\Sigma = \{S \subseteq V(G) : |S \cap A_i| = n - 1 \text{ for } 1 \leq i \leq r \text{ and } |S \cap B| = m_0\}$.

Let $G^* \supseteq G$ be defined as follows: $V(G^*) = V(G) \cup \{v_S : S \in \Sigma\}$. G^* induces G on $V(G)$. The v_S for $S \in \Sigma$ form an independent set, and the neighbors of v_S in $V(G)$ are the elements of S .

Example 12 For the simplest example, take $r = 1, n = 2, m = 1$. Then $G = K_2 + K_1$, and $G^* \simeq C_5$.

Lemma 8 G^* defined above has chromatic number $n + 1$.

Proof:

One can color G^* with $n + 1$ colors by first coloring G with n colors, and using the last color for all remaining vertices.

On the other hand, if G^* is colored with n colors, then all n colors occur in A_1 . Fix $b \in B$. For each color c fix $S(c) \in \Sigma$ so that $b \in S(c)$ and $S(c) \cap A_1$ consists of those vertices not of color c . Thus $v_{S(c)}$ must have color c , so b does not have color c . Therefore b cannot be colored. ■

Now we give additional examples of constraint families \mathcal{C} such that the algebraic closure is trivial in $\mathcal{E}_{\mathcal{C}}$ (i.e.. $\text{acl}(A) = A$), and $|\mathcal{C}| = 2$.

Proposition 2 For r, m, n with $n \geq m \geq 1$ and either $r \geq 2$ or $m \geq n - 1$, and for G^* defined as above, $\mathcal{H}(G^*) = \{G^*, K_{n+1}\}$. In particular $|\mathcal{H}(G^*)| = 2$.

Proof:

Evidently G^* does not contain K_{n+1} . It suffices now to prove that any proper homomorphic image of G^* does contain K_{n+1} .

Let $h : G^* \rightarrow H$ with $h(u) = v$, for some $u, v \in V(G^*), u \neq v$. Note that u, v are not adjacent. We consider cases.

Case 1. $u \in A_i$ for some i , and $v \in G$.

Take $S \in \Sigma$ with $S \cap A_i = A_i - \{u\}$, $v \in S$. Then the induced graph on $\{v_S\} \cup A_i$ is isomorphic to K_{n+1} in H .

Case 2. $u \in G$, for some i , and $v \notin G$.

Let $A^* = A_i$ or B be the component of G containing u . As $r \geq 2$ or $m \geq n - 1$, we may choose $A \subseteq G, A \simeq K_{n-1}$, with $A \cap A^* = \emptyset$, and v adjacent to all vertices of A . Take S with $u \in S, A \subseteq S$. Then in H , the induced graph on $\{v\} \cup \{A\} \cup \{v_S\}$ is isomorphic to K_{n+1} .

Case 3. $u, v \notin G$

Let $u = v_S, v = v_T$. Let A be a connected component of G such that $A \cap S \neq A \cap T$. Then in H , the induced graph on $\{u\} \cup A$ is isomorphic to K_{n+1} . ■

§8. Another Universal Graph.

The main result of this section is that for the graph $C = T_1 + T_2 + P_n$ consisting of two triangles T_1, T_2 with exactly one common vertex and a path of length n starting from a non-common vertex in one of these triangles, the theory T_C^* is \aleph_0 -categorical. Here we use the *ad hoc* notation $+$ for an almost disjoint sum with one (specific) pair of vertices identified. (We write $+_v$ when the common vertex v needs to be specified). This depends on an analysis of algebraic closure much of which is valid more generally and may be useful in the analysis of other candidates for membership in \mathcal{U}_0 (defined in the introduction).

We will assume throughout that \mathcal{C} is a finite set of finite connected graphs. Furthermore G denotes an \aleph_0 -saturated graph in $\mathcal{E}_{\mathcal{C}}$. We use the term “weak embedding” for the ordinary graph theoretic embedding (as opposed to a strict embedding, which is an isomorphism with an induced subgraph).

Definition 8

1. For $A \subseteq H \subseteq G$ with H finite, we say that H is free over A if there is an embedding of infinitely many copies of H in G over A , disjoint over A (this means that the intersection of any two copies is A).
2. For $A \subseteq H \subseteq G$ with H finite, $\text{cl}(A; H)$ is the union of A with all sets B such that:
 - 2.1 H is free over $A \cup B$;
 - 2.2 B is minimal subject to 2.1.
3. Let \mathcal{F} be a collection of pairs (A, H) of finite graphs with $A \subseteq H$. Then for any $X \subseteq G$, $\text{cl}(X; \mathcal{F})$ is the union of all sets of the form $\text{cl}(A; H)$ where $A \subseteq X, H \subseteq G$, and (A, H) is isomorphic to a pair in \mathcal{F} .
4. With \mathcal{F} as in 3, we say that \mathcal{F} is a base for acl if for all $X \subseteq G$ we have: $X = \text{cl}(X; \mathcal{F})$ if and only if $X = \text{acl}(X)$.
5. A graph H is solid if every induced 2-connected subgraph of H is complete.

We may now state the main results:

Proposition 3. For any pair A, H of finite graphs with $A \subseteq H \subseteq G$, we have $\text{cl}(A; H) \subseteq \text{acl}(A)$. Hence for any collection \mathcal{F} of pair (A, H) of finite graphs with $A \subseteq H$, and any $X \subseteq G$, $\text{cl}(X; \mathcal{F}) \subseteq \text{acl}(X)$.

Proposition 4. If \mathcal{F} is a finite set of pairs (A, H) of finite graphs with $A \subseteq H$, and $X \subseteq G$ is finite, then $\text{cl}(X; \mathcal{F})$ is finite.

Proposition 5. Let $\mathcal{F} = \{(A, H) : A \subseteq H \subseteq G \text{ and for some } C \in \mathcal{C}, H \text{ embeds weakly in } C \text{ as a proper subgraph of } C\}$. Then \mathcal{F} is a base for acl .

Proposition 6. If \mathcal{C} consists of solid graphs, and if \mathcal{F} is the collection of pairs $(\{a\}, H)$ for which $a \in H$, H embeds properly in some $C \in \mathcal{C}$, and $H - \{a\}$ is a connected component of $C - \{a\}$, then \mathcal{F} is a basis for acl . In particular:

$$(*) \quad \text{acl}(X) = \bigcup_{a \in X} \text{acl}(a)$$

for $X \subseteq G$.

We do not know exactly when the “unarity” condition $(*)$ holds; it might be useful to determine this. If we take the union of a collection of solid graphs and a collection closed under homomorphic image, then the same property holds since acl is unchanged. However, if $\mathcal{C} = \{C\}$ consists of a single forbidden subgraph, then $(*)$ is equivalent to the solidity of C .

Definition 9

The next statement requires a more delicate partial closure operation, for use with the particular graph $C = T_1 +_{u_2} T_2 +_{y_0} P_n$ referred to above. Let \mathcal{F} be the set of pairs $(\{a\}, P)$ for which P is a path of length at most n with a an endpoint. For $X \subseteq G$ let $\text{cl}_C^*(X)$ be the union of $\text{cl}(X; \mathcal{F})$ with

1. all sets of the form $\text{cl}(\{a\}; H)$ for which: $a \in X$; a lies in some copy of $T_1 + T_2$ with a not the common vertex of the two triangles; $H \simeq T + P$, the free amalgam of a triangle T with a path of length at most n ; and
2. the set of all points b lying in $\text{cl}(\{a\}; H)$ with $a \in X$, $a \in H \simeq T + P$, and either a, b belong to a triangle, or b lies in some copy of $T_1 + T_2 + P$ with b not the common vertex of the two triangles, and with P a path.

Proposition 7. Let $C = T_1 +_{u_2} T_2 +_{y_0} P_n$ be the graph referred to above, obtained by amalgamating two triangles T_1, T_2 , and a path P_n of length n , over two distinct points of T_2 . Then for $X \subseteq G$ and cl_C^* as defined above, if $X = \text{cl}_C^*(X)$ then $X = \text{acl}(X)$.

Proposition 8 Let $C = T_1 +_{u_2} T_2 +_{y_0} P_n$ be the graph referred to in the previous Proposition. Then the theory T_C^* is \aleph_0 -categorical, and thus there is a universal C -free graph.

Proof of Proposition 3:

We consider $A \cup B \subseteq H$ with H free over $A \cup B$, and with B minimal subject to this condition (in particular $A \cap B = \emptyset$). If $B \not\subseteq \text{acl}(A)$, then there are infinitely many copies (B_i, H_i) of (B, H) embedded as induced subgraphs of G , with the B_i distinct and the H_i free over B_i . Without loss of generality the B_i form a Δ -system with common part B_0 . As the B_i are disjoint over B_0 and each H_i is free over $A \cup B_i$, H is free over $A \cup B_0$. This contradicts the minimality of B . ■

Proof of Proposition 4:

It is easy to see that $\text{cl}(X; \mathcal{F})$ is a definable set, and as it is contained in $\text{acl}(X)$, and G is \aleph_0 -saturated, it is finite. For the definability it suffices to check the definability of “free over”; but we can replace the requirement of infinitely many disjoint copies of H by k disjoint copies, where $k = \max\{|C| : C \in \mathcal{C}\}$, since $G \in \mathcal{E}_{\mathcal{C}}$. ■

Proof of Proposition 5:

Let $\mathcal{F} =$

$$\{(A, H) : A \subseteq H \subseteq G \text{ and for some } C \in \mathcal{C}, H \text{ embeds weakly in } C \text{ as a proper subgraph of } C\}$$

Let $X \subseteq G$, and assume that $X = \text{cl}(X; \mathcal{F})$. We claim $X = \text{acl}(X)$. We may suppose that X is finitely generated.

Suppose $X \neq \text{acl}(X)$. Then as G is \aleph_0 -saturated, if we form $G(2) = G_1 +_X G_2$ with G_1, G_2 isomorphic to G over X , then $G(2) \notin \mathcal{G}_{\mathcal{C}}$, and thus there is a weak embedding $h : C \hookrightarrow G(2)$ for some $C \in \mathcal{C}$. Let $H_2 = h[C] \cap G_2$, and let H_1 be the image of H_2 in G_1 under the given isomorphism. As $X = \text{cl}(X; \mathcal{F})$ and the pair $(X \cap h[C], H_1)$ lies in \mathcal{F} , H_1 is free over $X \cap h[C]$, and hence can be embedded in G_1 disjoint from $G_1 \cap h[C]$ over $X \cap h[C]$. Defining h' to agree with h off $h^{-1}[G_2]$ and with this new embedding on $h^{-1}[G_2]$, we have an embedding of C into G_1 . As $G_1 \in \mathcal{G}_{\mathcal{C}}$ this is a contradiction. ■

The next two proofs will be somewhat similar to the foregoing, and very similar to one another.

Proof of Proposition 6:

\mathcal{C} consists of solid graphs and \mathcal{F} is the collection of pairs $(\{a\}; H)$ for which $a \in H$, H embeds properly in some $C \in \mathcal{C}$, so that $H - \{a\}$ is a connected component of $C - \{a\}$. The proof that follows will allow us to replace \mathcal{F} by a slightly more restricted family which will be defined below.

We take $X \subseteq G$ finitely generated (with respect to this closure operation) and we suppose that $X = \text{cl}(X; \mathcal{F})$ but $X \neq \text{acl}(X)$, so that after forming $G(2) = G_1 +_X G_2$ as in the previous argument, we have an embedding $h : C \hookrightarrow G(2)$ for some $C \in \mathcal{C}$. We associate to C the tree T whose vertices correspond to the 2-connected components of C , with edges between components which either meet or are connected by some edge of C . We will denote the vertices of T by t, t' and the like, and the component of C corresponding to a vertex t of T will be denoted C_t . Now pick an arbitrary vertex 0 of T , and take it as a root for T . Now T can be viewed as partially ordered set with minimum 0. For $t \in T$ let $T^t = \{t' : t' \geq t\}$ and let $C^t = \bigcup_{t' \in T^t} C_{t'}$.

We will replace the set \mathcal{F} considered above by the subset of pairs $(\{a\}; H)$ for which for some $t > 0$ either:

1. $a \in C_t$ and $H - \{a\}$ is a component of $C^t - \{a\}$; or
2. t is a successor of a node t^- , $a \in C_{t^-}$, and $H - \{a\}$ is C_t .

Note that in the first case, typically $H = C^t$; this holds for example if $|C_t| > 1$.

Of course, with this modification we still have $X = \text{cl}(X; \mathcal{F})$.


Now let t be maximal in T such that $h[C^t]$ does not lie in either factor G_1 or G_2 of $G(2)$. We may suppose that $h[C_t] \subseteq G_1$. It will suffice to replace $h : C \rightarrow G(2)$ on C^t by h' agreeing with h on $(C - C^t) \cup C_t$, so that $h'[C^t - C_t] \subseteq G_1$; repeating this operation eventually produces an embedding of C into G_1 , and a contradiction.

We may break this down two steps further. First, for each successor node t' of t for which $h[C^{t'}] \subseteq G_2$, it suffices to find an embedding $h' : C \rightarrow G(2)$ agreeing with h on $C - C^{t'}$ and taking $C^{t'}$ into G_1 . For the second step, first choose a vertex $a \in X$ as follows: if $C_t \cap C_{t'} \neq \emptyset$, let a be the unique vertex common to both components. Otherwise, take a pair of vertices $u \in C_t$ and $v \in C_{t'}$ with u, v adjacent in C , and let $a = u$ if this is in X , and $a = v$ otherwise. With these choices, $a \in X$. Now we adjust h on $\{a\} \cup C^{t'}$ by making separate adjustments on each subgraph H containing a such that $H - \{a\}$ is a connected component of $C^{t'} - \{a\}$.

At this point the pair $(\{a\}; H)$ under consideration is one of the pairs which we have put in \mathcal{F} . As $X = \text{cl}(X; \mathcal{F})$, we can embed H freely into G_1 over a and arrive at the desired modification of h . Iterating this construction over all such components and all such nodes t' , and then over all suitable t , we will reach a contradiction. \blacksquare

Proof of Proposition 7:

We now deal with the particular case $C = T_1 +_{u_2} T_2 +_{y_0} P_n$, whose vertices we label as follows:



We follow exactly the same line as in the previous proof. Now the tree T is a path of length $n + 2$ whose first node corresponds to the first triangle T_1 ; take this node as a root and use the corresponding set \mathcal{F} of pairs:

$$\begin{aligned} &(\{a\}; P) \text{ with } a \text{ the initial vertex of a path } P \text{ of length at most } n \\ &(\{u_2\}, T_2 +_{y_0} P_n) \end{aligned}$$

This is almost what we want, except that the second possibility is somewhat more generous than we wish to allow. Accordingly, we will now consider the corresponding part of the previous argument more carefully. This occurs when the vertex t is the root and t' corresponds to the triangle T_2 , $a = u_2$, and we wish to embed $H = T_2 +_{y_0} P_n$ into G_1 over a disjoint from the image of T_1 . That is, we have an embedding $h : C \rightarrow G(2)$ with $h[T_1] \subseteq G_1$; $h[H] \subseteq G_2$, (so $h(u_2) \in X$) and we assume toward a contradiction that any embedding of H into G_1 meets $h[T_1]$. Let $b = h(u_3)$ and $c = h(y_0)$.

If $b, c \in X$ then it suffices to embed P_n into G_1 correctly, and this we have already dealt with. If neither b nor c lies in $\text{cl}(\{a\}; H)$ then T_2 can be embedded freely in G_2 over a , which produces a copy of C in G_2 since we already have $T_2 +_{y_0} P_n$ embedded in G_2 .

Thus we are left with the cases in which b or c lies in $\text{cl}(\{a\}; H)$, and in particular lies in X , and the other vertex is not in X .

Suppose $c \in \text{cl}(\{a\}; H)$ and $b \notin X$. Let $B = \text{cl}(\{a\}; H)$. Then we can embed B freely in G_1 over X , and then continue to embed P_n freely in G_1 over y_0 . This produces the desired embedding of C in G_1 (since “freely” means: without any undesirable identifications).

Finally, suppose $b \in \text{cl}(\{a\}; H)$ and $c \notin X$. In particular c is not in $\text{cl}(\{a\}; H)$ and hence there are infinitely many triangles containing a, b . Let $B = \text{cl}(\{a, b\}; H)$. We will show that $B \subseteq X$. Take $u \in B - \{a, b\}$, and set $B' = B - \{u\}$. Let G' be the free amalgam of G_2 with a large number of copies of H over B' . Then C embeds in G' .

Suppose this embedding involves a triangle $T = \{a, b, c'\}$ lying in one of the copies of H . Then there is a triangle T_0 meeting $\{a, b\}$ in a single vertex. If this triangle contains c , then choosing $c^* \in G_2$ not in $T_0 \cup h[P_n]$, so that $\{a, b, c^*\}$ lie on a triangle, we get an embedding of C into G_2 . So T_0 meets $\{a, b, c\}$ in one vertex. If T_0 contains u then by definition $u \in cl_C^*(X) = X$. So suppose it does not contain u . As $T \cup T_0$ is part of a copy of C embedded in G' , this embedding also involves a path P of length n attached to T or T_0 , and not at their common point. If the path is attached to T_0 , then replacing c by a point c^* for which a, b, c^* forms a triangle and c^* lies off $T_0 \cup P$, again C embeds in G_2 , a contradiction. So P is attached to T . P is broken into various connected components by its intersection with B' . We will alter the embedding so that P becomes a path of length n attached to c and otherwise disjoint from $T_0 \cup \{a, b, c\}$. Those segments which lie in G_2 may be left as they are. The remainder lie in copies of H , are attached at one or two points of $B' \cup \{c'\}$, and correspond to segments in G_2 which are either free over B' , or contain the point u . As u does not lie on T_0 , if a segment corresponding to one containing u occurs, it may be replaced by two segments in G_2 joined at u , and free over B . Thus by choosing the embedding of P carefully, one may embed C in G_2 , a contradiction.

Therefore in our original embedding of C into G' , the copy of $T_1 + T_2$ embeds in G_2 and part of the path P is embedded in various copies of H amalgamated over B' . Again we can alter most of the embedding of P to go into G_2 , apart from segments which correspond to segments in G_2 lying between two successive points of B' , with the vertex u on the segment. If any such segment actually occurs, it means that in G_2 , u lies on some graph of the form $T_1 + T_2 + P$ with P a path. Thus again $u \in cl_C^*(\{a\})$. ■

For the next proof we will require an auxiliary result which will be seen to contain useful information about algebraic closure in the case at hand. Let $C = T_1 +_{u_2} T_2 +_{y_0} P_n$ be the graph referred to in Proposition 8. In particular, n is fixed.

Lemma 9 *Let G be a graph, u a vertex of G , and suppose that there are two disjoint paths of length $5n$ originating at u , as well as an embedding of some subgraph H of C of the form $C = T_1 +_{u_2} T_2 +_{y_0} P_k$ with $0 \leq k \leq n$, embedded with u as the terminal vertex of P_k . Then C embeds in G .*

Proof:

Let v be the vertex in G corresponding to the vertex u_2 of H . Let P be one of the two given paths, which does not contain v . Then P is broken into at most 5 connected components by its intersection with the vertices of $T_1 +_{u_2} T_2$ (as embedded in G), and one of these components has length at least n . Thus, if this intersection is nonempty, then C embeds in G .

Suppose P is disjoint from the image of $T_1 +_{u_2} T_2$ in G . Let y be the first vertex of the path P_k (starting from the vertex y_0 in T_2) which corresponds under the embedding to a vertex of P . Then on removal of y from P , one of the components has length at least n , and hence we again have an embedding of C into G . ■

Proof of Proposition 8:

We now wish to show that for X finite, $\text{acl}(X)$ is finite. We define inductively: $X_0 = X$, $X_{i+1} = cl_C^*(X_i)$, and we need to show that this process terminates. Suppose in fact that it goes on for k stages with k substantially larger than $10n$. Define a sequence of points $a_i \in X_i \setminus X_{i-1}$ for $i < k$ by downward induction so that $a_i \in cl_X^*(a_{i-1})$ for all i . The point a_{k-1} is selected arbitrarily, and given a_i , as it lies in X_i it lies in $cl_C^*(a_{i-1})$ for some $a_{i-1} \in X_{i-1}$, and this element lies outside X_{i-2} since $cl_C^*(a_{i-1})$ is not contained in X_{i-1} .

We claim:

(*) The elements a_i ($i < k$) can be selected so that they lie on a path of length at least $k - 1$.

Again, proceed by downward induction, building up a finite path Q_i with endpoint a_i as we go along in such a way that $Q_i \cap X_i = \{a_i\}$.

Suppose $b = a_i$ has been chosen and pick some $a \in X_{i-1}$ so that $b \notin cl_C^*(a)$. Suppose first that $b \in cl(\{a\}; P)$ for some path P , with a an endpoint of P . Let $B = cl(\{a\}; P)$ and let a' be the vertex of $B \cap X_{i-1}$ on the segment from a to b which is closest to b ; possibly $a' = a$. Let P' be the subpath of P with initial vertex a' , passing through b , and $B' = B \cap P'$. We know that P' is free over B' and easily B' is minimal with this property. Thus $B' \subseteq cl(\{a\}; P')$ and we may take $a_{i-1} = a'$. The remaining elements on the segment $(a', b) \cap B'$ are outside X_i and as $[a', b]$ is free over its intersection with B' , we may attach to Q_i a path from a_{i-1} to a_i which meets X_{i-1} only in a_{i-1} , and meets Q_i only in a_i ; this produces the desired path Q_{i-1} .

Now suppose that $b \in cl(\{a\}; H)$ with $H \simeq T + P_n$, $a \in T$, $a \notin P_n$. We can proceed in more or less the same way. If b is a vertex of T we can just take $a_{i-1} = a$ and adjoin the edge (a, b) to the path Q_i . Otherwise b lies on the path P_n . We consider $P = [a, b]$, the shortest path from a to b in H ; this meets T in two points. Let $B = cl(\{a\}; H) \cap P$ and let a' be the vertex of $B \cap X_{i-1}$ which is closest to b ; possibly $a' = a$. Let P' be the subpath of P with initial vertex a' , passing through b . If $a' \neq a$ we claim that $b \in cl(\{a'\}; P')$. This is seen as in the previous case. Furthermore P' is free over $B \cap P'$ so we may connect a' to b by a path meeting X_{i-1} in a alone, and meeting Q_i in b alone.

Thus we have (*), and in particular if $b_i = a_{5n+i}$ and i is not too large, we have two disjoint paths from b_i of length $5n$ contained in G . We claim that in this case every path of length n originating at b_i is free over b_i . If this fails, then as $G \in \mathcal{E}_C$ we must have a subgraph of C consisting of the two triangles and some initial segment I (possibly of length 0) of the path P_n , embedded with b_i as the terminal point of I . This violates the previous lemma. It follows that for $a = b_i$ and $b = b_{i+1}$ the relationship $b \in cl_C^*(a)$ is realized in the following way: $b \in cl(\{a\}; H)$ with $H \simeq T + P$ (a triangle amalgamated with a path of length n), and either a lies in some copy of $T_1 + T_2$ with a not the central vertex, contradicting the previous lemma, or b lies in some copy of $T_1 + T_2 + P$, with P a path and with b not the central vertex, again contradicting the previous lemma if $i + 1$ is not too large, or finally: a and b lie on a common triangle.

Thus we may assume that we have 4 consecutive points b_i ($i = 1, 2, 3, 4$) such that for $i = 1, 2, 3$ the pair (b_i, b_{i+1}) lies on a triangle (b_i, b_{i+1}, c_i) . This gives an embedding of $T_1 + T_2$ into $\{b_i, c_j : 1 \leq i \leq 4, 1 \leq j \leq 3\}$: if $c_1 = c_3$ use the first and third triangles, while otherwise these two triangles are disjoint and hence the second one meets at least one of them in a single vertex. This again violates the previous lemma since at least one of the b_i occurs as a noncentral point in the embedded copy of $T_1 + T_2$.

This contradiction completes the proof.

§9. Paths.

The existence of a universal P -free graph, when $P = P_k$ is a finite path of length k (and thus of order $k + 1$) is established in [KMP]. The analysis given there yields good structural information and allows further generalization, for example to categories of vertex colored graphs, which will be of further use even in the case of graphs. However it does not give realistic control over the sizes of algebraic closures. Writing $c_k(n)$ for

$\max |\text{acl}(A)|$, where A varies over sets of n vertices in graphs G belonging to \mathcal{E}_{P_k} , we would get an estimate of $c_k(n)$ of the form a tower of exponentials of height about k , using [KMP]. However the analysis of §8 yields:

$$c_k(n) = c_k(1) \cdot n$$

which is already fairly good, leaving open only the question of the growth rate of $c_k(1)$ as a function of k , which turns out to be an intriguing question. Consideration of circuits of length k , or, for that matter, any hamiltonian graphs of order k , yields:

$$c_k(1) \geq k$$

and for low values of k one may check $c_k(1) = k$. In fact the following is open:

Problem. Is $c_k(1)$ equal to k for all k ?

In the remainder of this section we will prove:

Proposition 9 $c_k(1) < k^{3k^2}$.

That is, we reduce a tower of exponentials to a single exponential, but fall far short of the linear bound which may hold. This result requires a closer and more concrete analysis of the operation of algebraic closure, which begins by simply following through on the analysis given in §8 more generally.

On the basis of Proposition 6 we can describe the algebraic closure operation in \mathcal{E}_{P_k} as follows. Let \mathcal{F} be the collection of pairs $(\{a\}, Q)$ where Q is a path of length at most $k-1$ and a is an endpoint of Q . For $G \in \mathcal{E}_{P_k}$ and $v \in V(G)$, define inductively:

$$\text{cl}_0(v) = \{v\}; \quad \text{cl}_{n+1}(v) = \text{cl}(\text{cl}_n(v); \mathcal{F}) - \text{cl}_n(v)$$

if $\bigcup_{i=0}^{k-1} \text{cl}_i(v) \neq \text{acl}(v)$, one produces a contradiction by constructing a path of length k by downward induction, beginning with some $u \in \text{cl}_k(v)$. Thus if we have a uniform estimate of the form $|\text{cl}_1(v)| \leq N$ holding in \mathcal{E}_{P_k} (all v), then correspondingly

$$\sup |\text{acl}(v)| \leq \sum_{i \leq k-1} N^i < N^k, \text{ (the supremum is over all } v \in G \in \mathcal{E}_{P_k})$$

We will get such an estimate with $N = k^{3k}$.

Definition 10 Let $G \in \mathcal{E}_{P_k}$.

1. For $u, v \in V(G)$, set $\omega(u, v) =$

$$\sup \{m : \text{There are infinitely many paths of length } m \text{ connecting } u \text{ and } v \text{ in } G, \\ \text{disjoint except for their endpoints}\}$$

This supremum is taken to be 0 if there is no such m . However when u and v are adjacent the condition is considered to hold with $m = 1$, in a degenerate form.

2. For $u \in V(G)$, set $\omega(u, \infty) =$

$$\sup \{m : \text{There are infinitely many paths of length } m \text{ in } G \text{ with } u \text{ as an endpoint,} \\ \text{disjoint apart from } u\}$$

Lemma 10. Let $G \in \mathcal{E}_{P_k}$, $v \in V(G)$, and suppose that Q is a path in G originating at v , while $B \subseteq Q - \{v\}$ is minimal such that Q is free over $B \cup \{v\}$. Write $B \cup \{v\}$ as a sequence (v_0, v_1, \dots, v_l) in order along Q , beginning with $v_0 = v$. Then $\omega(v_i, v_{i+1}) \geq 1$ for all $i < l$ and:

1. $\omega(v_i, v_j) < \sum_{i \leq r < j} \omega(v_r, v_{r+1})$ for $i \leq j - 2$; and
2. $\omega(v_i, \infty) < \sum_{i \leq r < l} \omega(v_r, v_{r+1}) + \omega(v_l, \infty)$ for $i < l$.

Proof:

Since Q is free over $B \cup \{v\}$, we have $\omega(v_i, v_{i+1}) \geq 1$.

Condition (1) follows easily from the assumption that Q is not free over $(B \cup \{v\}) \setminus \{v_r : i < r < j\}$ for $i \leq j - 2$, and condition (2) follows from the assumption that Q is not free over $(B \cup \{v\}) \setminus \{v_r : r > i\}$. ■

Definition 11 Let $G \in \mathcal{E}_{P_k}$ and let $\mathbf{v} = (v_0, \dots, v_l)$ be a sequence of vertices of G .

1. \mathbf{v} is a chain if $l \geq 2$, $\omega(v_i, v_{i+1}) \geq 1$ for all i , and $\omega(v_i, v_j) < \sum_{i \leq r < j} \omega(v_r, v_{r+1})$ for $i \leq j - 2$.
2. Similarly, the formal sequence $(v_0, \dots, v_l, \infty)$ is called an open chain if $l \geq 1$ (so the length is at least 2) and if it satisfies the same formal conditions, using $\omega(v_i, \infty)$ where called for.
3. The virtual length of a chain \mathbf{v} is $\sum_r \omega(v_r, v_{r+1})$, and the virtual length of an open chain is defined similarly. We write $\lambda(\mathbf{v})$ for the virtual length of \mathbf{v} .

In the proof of the next lemma we will need a result of Erdős and Gallai:

Fact [EG, Theorem 2.6] let H be a graph with n vertices and e edges, in which there is no path containing l edges ($l \geq 1$). Then $e \leq \frac{n(l-1)}{2}$.

Lemma 11. Let $G \in \mathcal{E}_{P_k}$, $v \in V(G)$, and $A \subseteq \text{cl}_1(v)$. Then there is a set $A' \supseteq A$, with $|A' - A| < (k^3/4)|A|$, such that for any chain or open chain \mathbf{v} whose endpoints lie in $A \cup \{\infty\}$, if \mathbf{v} is not contained in $A \cup \{\infty\}$ then it meets $A' \setminus A$.

Proof:

For each $a \in A$ choose one path originating at a , of maximum length, and let \mathcal{A}_1 be the set of paths chosen. Let \mathcal{A}_2 be a maximal collection of chains whose endpoints lie in A , and which are otherwise disjoint both from each other and from the paths in \mathcal{A}_1 . We will take $A' = \bigcup(\mathcal{A}_1 \cup \mathcal{A}_2)$. There are a number of points to be verified. We will begin by verifying that A' has the desired property, then estimate its size.

Consider first a chain with endpoints in A , not wholly contained in A . We may suppose then that only its endpoints lie in A . By the choice of \mathcal{A}_2 , if this chain does not meet any path in \mathcal{A}_1 in one of its interior points, then it meets one of the chains in \mathcal{A}_2 .

Now consider an open chain \mathbf{v} originating at a vertex a of A , and not wholly contained in $A \cup \{\infty\}$. Then we may suppose that it meets A only at a , as otherwise we would replace it either by a shorter open chain, or by a chain with endpoints in A . Let $L = \lambda(\mathbf{v})$. Then $\omega(a, \infty) < L$, so there is a path (an “obstruction”) of length at least $k - L$ with a as an endpoint. Therefore there is such a path in \mathcal{A}_1 , and it is easy to see that \mathbf{v} meets that path at an interior point, as otherwise one constructs a path of length k in the ambient graph.

For cardinality estimates it will be convenient to take $k \geq 3$, as we may. We have $|\mathcal{A}_1| \leq |A|$ and to complete the analysis we will show:

$$|\mathcal{A}_2| \leq \frac{(k-1)(k-1)}{4}|A|,$$

from which our claim follows easily.

To make this estimate, we will estimate separately the number of chains in \mathcal{A}_2 connecting two specified vertices, and the number of pairs of vertices having such a connection.

We begin with the latter point. Consider the graph Γ whose vertex set is A , and with edges between pairs of vertices joined by one of the chains in \mathcal{A}_2 . Our claim is:

$$(*) \quad e(\Gamma) \leq \frac{k-1}{4}|A|$$

We claim that Γ contains no path of length $\lceil k/2 \rceil$; this property then implies $(*)$ by the result of Erdős and Gallai [EG] which was quoted above. If Γ had a path of length $\lceil k/2 \rceil$, that is to say a sequence of at least $k/2$ chains which are disjoint except at their endpoints, then these chains fit together to form a path of length at least k (extending the chains “freely” to their virtual lengths). Note that each chain has virtual length at least 2 by definition, and the chains may be extended so that the added vertices are distinct from each other and any vertices previously considered.

Now we consider the number $\mu(a, b)$ of chains in \mathcal{A}_2 which connect two fixed vertices $a, b \in A$. We claim $\mu(a, b) \leq k - 1$. To see this, let $l = \omega(a, b)$. Then for each chain $\mathbf{v} \in \mathcal{A}_2$ which connects a and b , we have $\lambda(\mathbf{v}) \geq l + 1$. Let G_1 be the free amalgam of the ambient graph G with infinitely many additional paths of length $l + 1$ connecting a to b . If $G_1 \in \mathcal{G}_{P_k}$ then as $G \in \mathcal{E}_{P_k}$ and $\omega(a, b) \geq l + 1$ in G_1 , we find $\omega(a, b) \geq l + 1$ in G , a contradiction.

Thus G_1 contains a path P of length k . $P \setminus \{a, b\}$ consists of at most 3 segments, each of which lies either wholly in G or in one of the additional paths of length $l + 1$ adjoined to form G_1 . Assuming $\mu(a, b) > 0$, there is at least one path in G of length $l + 1$ joining a and b . Therefore we may suppose that P is chosen so that $P' = P \cap G$ contains at least one of the segments of $P \setminus \{a, b\}$. In particular $P \setminus P'$ consists of at most two segments. We now count separately the chains $\mathbf{v} \in \mathcal{A}_2$ connecting a and b which meet $P' \setminus \{a, b\}$, and those which do not. P' contains at most $k - 2$ vertices and hence meets at most $k - 2$ of the chains in \mathcal{A}_2 . Furthermore P' is disjoint from at most one chain in \mathcal{A}_2 which links a and b , as two such chains could be extended freely to give two disjoint paths of length $l + 1$ joining a and b , into which the segments of $P \setminus P'$ could be copied, thereby embedding P_k in G . Thus there are at most $1 + (k - 2) = k - 1$ chains in \mathcal{A}_2 linking a and b . This completes our estimate. ■

Corollary For $G \in \mathcal{E}_{P_k}$, $v \in G$, we have $|\text{cl}_1(v)| < k^{3k}$.

Proof:

Let $A_0 = \{v\}$ and define inductively $A_{i+1} = A'_i$ in the sense of Lemma 11 (this is not canonical, of course). In other words, choose A_{i+1} satisfying:

$$A_i \subseteq A_{i+1}, \quad |A_{i+1} \setminus A_i| < \frac{k^3}{4}|A_i|,$$

so that any chain or open chain \mathbf{v} whose endpoints lie in $A_i \cup \{\infty\}$ is either contained in $A_i \cup \{\infty\}$, or meets $A_{i+1} \setminus A_i$.

Then $|A_{i+1}| \leq k^3|A_i|$ and hence $|A_i| \leq k^{3i}$ for all i . On the other hand $\text{cl}_1(v) \subseteq A_k$ since each open chain \mathbf{v} originating at v will meet $A_{i+1} \setminus A_i$, as long as it is not contained in A_i ; and \mathbf{v} has at most k vertices. ■

§10. *More examples of universal graphs.*

We give two more examples of constraints allowing universal graphs. These are less complex than the family treated in §8, and may allow some further elaboration.

Consider first the constraint $C = K_n + P_k$ consisting of a complete graph with an attached path. It can be shown using either the methods of §8 or those of [KMP] that \mathcal{E}_C is \aleph_0 -categorical. Every connected component G_0 of a graph $G \in \mathcal{E}_C$ either omits K_n and belongs to \mathcal{E}_{K_n} , or contains a copy K of K_n , in which case the connected components of $G_0 \setminus K$ omit P_{2k} , and a structure theory for these can be given in the spirit of [KMP]. Alternatively, following the argument of §8, one finds that a vertex lying on a sufficiently long path has trivial algebraic closure.

Our second example is a slight generalization of the bow-tie, namely $C = K_n + K_3$, a complete graph attached to a triangle. A detailed analysis of the algebraic closure operator in this case will yield:

$$|\text{acl}(A)| \leq (n+1)|A|$$

for $A \subseteq G \in \mathcal{E}_C$. We will now give the details.

Definition 12 *Let $G \in \mathcal{E}_C$.*

1. *For $A \subseteq V(G)$, A is special if G induces a complete graph on A , and one of the following occurs:*
 - a. *$|A| = n$ and there is no $B \subseteq V(G)$ of order n such that $B \neq A$, $B \cap A \neq \emptyset$, and G induces a complete graph on B ; or*
 - b. *$|A| = n - 1$ and there are at least two vertices of G adjacent to all vertices of A .*
3. *For $a \in V(G)$ set $a^* =$*

$$\bigcup \{B : B \text{ is special, and the graph induced on } \{a\} \cup B \text{ is complete}\}$$

Remark

Let $G \in \mathcal{E}_C$, $A \subseteq V(G)$ of order n , and suppose that G induces a complete graph on A . Then A contains a special subset.

Our objective is to show that for $G \in \mathcal{E}_C$ and $a \in V(G)$, we have $\text{acl}(a) = \{a\} \cup a^*$, and $|\{a\} \cup a^*| \leq n+1$, so that by Proposition 6 of §8 we may conclude $|\text{acl}(A)| \leq (n+1)|A|$ for all A , and in particular this constraint allows a (canonical) universal countable graph.

Lemma 12 *Let $G \in \mathcal{E}_C$.*

1. *If $A, B \subseteq V(G)$ are of order n , and G induces a complete graph on each, then either $A = B$, or $A \cap B = \emptyset$, or $|A \cap B| = n - 1$.*
2. *If $A, B \subseteq V(G)$ are special and $A \cap B \neq \emptyset$, then either $A = B$ or $A \cup B$ is contained in a set of vertices of order $n + 1$ on which G induces a complete graph.*

Proof:

(1) holds by inspection.

For (2), if $|A| = n$ the claim holds by definition, so suppose that $|A| = |B| = n - 1$ and $A \neq B$. Let $A^+ = A \cup \{a\}$ and $B^+ = B \cup \{b\}$ be two sets of vertices of order n on which G induces complete graphs. We may suppose that they are chosen so that $a \notin B$ and $b \notin A$. This forces $a = b$ as $|A^+ \cap B^+| = n - 1$. Now let

$B^* = B \cup \{b'\}$ be another choice for B^+ , so $b' \neq b$. If $b' \notin A$ then G induces C on $(B \cup \{b'\}) \cup (\{a\} \cup A \setminus B)$, a contradiction. So $A \setminus B = \{b'\}$, which means that G induces K_n on $A \cup B$, and induces K_{n+1} on $\{a\} \cup A \cup B$. ■

Corollary For $G \in \mathcal{E}_C$, and $a \in V(G)$, we have $|\{a\} \cup a^*| \leq n + 1$.

Lemma 13 For $G \in \mathcal{E}_C$, and $A \subseteq V(G)$, the following are equivalent:

1. $A = \text{acl}(A)$;
2. $a^* \subseteq A$ for $a \in A$.

Proof:

(1 \Rightarrow 2): Assume $A = \text{acl}(A)$, $a \in A$, and $B \subseteq a^*$ is special, with G inducing a complete graph on $\{a\} \cup B$. If $|B| = n$ this implies that $a \in B$ and B is the unique such set containing a , hence belongs to $\text{acl}(a) \subseteq A$. If on the other hand $|B| = n - 1$ then either B is unique, or else $B \cup \{a\}$ is contained in a unique complete subgraph of G of order $n + 1$, by part (2) of the preceding lemma. In either case $B \subseteq \text{acl}(a) \subseteq A$.

(2 \Rightarrow 1): We suppose $a^* \subseteq A$ for $a \in A$, but A is not algebraically closed, and hence C embeds into the free amalgam $G_1 +_A G_2$ of two copies of G amalgamated over A . Let C^* be the image of C . As $C^* \setminus A$ is disconnected, there is a point $a \in C^* \cap A$ which lies on a complete graph K of order n and a triangle T intersecting at a ; we may take K to lie in G_1 , and T to lie in G_2 .

In particular $V(K)$ contains a special set B , and then $B \subseteq a^* \subseteq A$. As G_2 is C -free, this forces $|\{a\} \cup B| = n - 1$, that is: $|B| = n - 1$ and $a \in B$. Now let T_1 be the triangle in G_1 which corresponds to T in G_2 (via some isomorphism over A). Then $V(T) \cap B = \{a\}$ but T_1 must have another vertex in common with K , as G_1 is C -free, and thus $T_1 \supseteq V(K) \setminus B$ (which consists of a single vertex). At the same time, as B is special, there is another complete graph K' of order n in G_1 which contains B , and by the same token the triangle T_1 contains $V(K') \setminus B$; so as T_1 is complete, it follows that the graph induced on $V(K) \cup V(K')$ is complete of order $n + 1$. But then $V(K) \subseteq a^* \subseteq A$ as all its subsets of order $n - 1$ are special, and again $C^* \subseteq G_2$, a contradiction. ■

Thus as indicated above, we find $|\text{acl}(A)| \leq (n + 1)|A|$, so \mathcal{E}_C is \aleph_0 -categorical and there is a universal C -free graph. The case of a bouquet of two complete graphs, each of order at least 4, has not been investigated and may well succumb to a similar analysis. In any case, we believe that it should now be clear that the classification of the class \mathcal{U}_0 , described in the introduction, is within reach, albeit this would involve some rather substantial computations in positive cases and some additional concrete constructions to cover the negative cases. We emphasize that while the details would no doubt be tedious, the result would be a reasonably well-founded conjecture as to the general solution of the problem of the existence of a countable universal graph, for the class of graphs specified by prescribing any (single) finite connected forbidden subgraph.

References

- [Ba] P. Bacsich, *The strong amalgamation property*, Colloquium Mathematics, **33** (1975), 13-23.
- [BM] J. A. Bondy and U. S. R. Murty, **Graph Theory with Applications**, North-Holland, 1979.
- [CK] C. C. Chang and H. J. Keisler, **Model Theory**, North-Holland, 3rd ed., 1990.

- [ChK] G. Cherlin and P. Komjáth, *There is no universal countable pentagon-free graph*, J. Graph Theory, **18** (1994), 337-341.
- [CS1] G. Cherlin and N. Shi, *Graphs omitting sums of complete graphs*, J. Graph Theory, **24** (1997), 237-247.
- [CS2] ———, *Graphs omitting a finite set of cycles*, J. Graph Theory, **21** (1996), 351-355.
- [CST] G. Cherlin, N. Shi, and L. Tallgren, *Graphs omitting a bushy tree*, J. Graph Theory, **26** (1997), 203-210.
- [EG] P. Erdős and T. Gallai, *On maximal paths and circuits of graphs*, Acta Math., Acad. Sci. Hung. **10** (1959), 337-356.
- [FK1] Z. Füredi and P. Komjáth, *On the existence of countable universal graphs*, J. Graph Theory, **25** (1997), 53-58.
- [FK2] ———, *Nonexistence of universal graphs without some trees*, Combinatorica, **17** (1997), 163-171.
- [GK] M. Goldstern and M. Kojman, *Universal arrow free graphs*, Acta. Math. Hungary **73** (1996), 319-326.
- [HW] J. Hirschfeld and W. H. Wheeler, **Forcing, Arithmetic, Division Rings**, Springer-Verlag, 1975.
- [Ko] P. Komjath, *Some remarks on universal graphs*, Discrete Math., to appear.
- [KMP] P. Komjáth, A. Mekler and J. Path, *Some universal graphs*, Israel J. Math. **64** (1988), 158-168.
- [KP1] P. Komjáth and J. Pach, *Universal graphs without large bipartite subgraphs*, Mathematika, **31** (1984), 282-290.
- [KP2] ———, *Universal elements and the complexity of certain classes of infinite graphs*, Discrete Math. **95** (1991), 255-270.
- [Pa] J. Pach, *A problem of Ulam on planar graphs*, Eur. J. Comb. **2** (1981), 357-361.
- [Ra] R. Rado, *Universal graphs and universal functions*, Acta Arith. **9** (1964), 331-340.